Learning and Expectations in Dynamic Spatial Economies*

Jingting Fan
Cheung Kong Graduate School of Business

Sungwan Hong
University of Pittsburgh

Fernando Parro
University of Rochester and NBER

September 2025

Abstract

The impact of shocks in dynamic spatial environments depends on how forward-looking agents anticipate the path of future fundamentals. We develop a framework that incorporates flexible beliefs about future fundamentals into a general class of dynamic spatial models, allowing for beliefs that are evolving, uncertain, and heterogeneous across agents. This framework provides a tractable methodology to quantify the consequences of both ex-ante and ex-post shocks. We apply it to two settings: an ex-post evaluation of the China productivity shock on the U.S. economy, and an ex-ante study of the economic impacts of climate change. In both cases, we study the impact of deviating from perfect foresight on key economic outcomes.

*For helpful comments and discussion, we thank Matilde Bombardini, Lorenzo Caliendo, Javier Cravino, Jonathan Eaton, Andrei Levchenko, Yuhei Miyauchi, Eduardo Morales, Nitya Pandalai-Nayar, Stephen Redding, Andrés Rodríguez-Clare, Natalia Ramondo, Esteban Rossi-Hansberg, Sebastian Sotelo, James Tybout, Jonathan Vogel, Michael Waugh and seminar and conference participants at Berkeley, Boston University, Chicago, Chinese University of Hong Kong, Emory, Georgetown, Michigan, Hong Kong University, City University of Hong Kong, Iowa State, NYU, NYU Abu Dhabi, Princeton, Penn State, Rochester, St. Louis Fed, Temple University, University of Florida, University of Houston, University of Pittsburgh, University of Washington Seattle, University at Buffalo SUNY, Yale, SED Cartagena, and NBER Summer Institute ITM. Correspondence by e-mail: jtfan@ckgsb.edu.cn; suh90@pitt.edu; fernando.parro@rochester.edu.

1 Introduction

In dynamic spatial economies, the impact of shocks to fundamentals, such as location-specific productivity, depends crucially on how forward-looking agents anticipate the path of these fundamentals. Although a frequent assumption in the spatial literature is that agents have perfect foresight, in reality, beliefs can diverge from perfect foresight in various ways. For example, agents may not know the process governing the path of fundamentals, forcing them to form beliefs based on history. Even when they understand the process, if the process is stochastic, agents face uncertainty and their expectations will generally differ from ex-post realizations.

Incorporating departures from perfect foresight into dynamic spatial economies presents two challenges. The first challenge lies in recovering counterfactuals when the data are generated by agents without perfect foresight. In many applications, researchers observe the actual outcome and use it to discipline models for counterfactual analysis of fundamental shocks. Consider, for instance, estimating the impact of China's past growth on U.S. employment. The observed employment evolution reflects the decisions of agents who may hold evolving and potentially inaccurate beliefs about the extent of China's rise. Ignoring this possibility and attributing a slow reallocation out of adversely affected regions solely to, for example, migration frictions, may result in biased parameter estimates and misleading counterfactual predictions. The second challenge lies in incorporating uncertainty about the path of fundamentals into dynamic models with many locations and sectors, which leads to a large state space and many sources of uncertainty, posing tractability challenges.

This paper addresses these challenges in a general class of dynamic spatial models. Our model, presented in Sections 2.1 to 2.4, builds on the recent spatial literature by incorporating forward-looking migration decisions and a general equilibrium trade structure. Without loss of generality, we allow productivity to follow generic *location-specific* stochastic processes, which can be non-stationary with arbitrary historical dependence, while other fundamentals are time-varying but deterministic. Agents make forward-looking decisions based on their beliefs about these processes. Crucially, beliefs can also be time-varying and history-dependent in flexible ways, accommodating different belief formation mechanisms.

As an intermediate step in addressing these challenges, we solve the model by approximating the stochastic equilibrium around a deterministic transition path. While closely related to approximation techniques around steady states that are standard in macroeconomics, our approach is tailored to accommodate the flexible, history-dependent beliefs and

time-varying fundamentals central to our analysis. We build on this approximate solution to tackle both challenges. For expositional clarity, we first leverage the first-order approximation to explain our method for recovering counterfactuals from the observed outcomes in settings without perfect foresight. We then show how to incorporate the impacts of uncertainty, which requires a second-order approximation.

Section 2.5 addresses the first challenge. Consider a researcher who observes outcomes shaped by agents with imperfect foresight and seeks to uncover outcomes under a counterfactual path of fundamental productivity. This researcher must first infer parameters that shape dynamic decisions, such as migration costs. However, because agents base their decisions on expected future productivity, not its ex-post realizations, the standard approach in the literature—using realized outcomes to infer parameters via model inversion—is invalid. We develop a methodology for conducting counterfactuals in such settings.

Our methodology requires data on agents' beliefs at different times about the path of fundamental productivity.[1] We combine this belief data with two key theoretical observations. First, conditional on agents' beliefs at time t, the sequence of *expected future endogenous outcomes* (e.g., wages, prices, migration flows)—which we term the 'expected path from t'—satisfies the agents' forward-looking equilibrium conditions at t (up to the first order). This expected path therefore encodes information about migration costs and other structural parameters. Second, while each of the expected paths is unobserved, the difference between, say, the expected path from t and the expected path from a later period $t+k$, is determined entirely by the belief errors (i.e., the discrepancies between realized and expected productivity) accumulated between t and $t+k$. This observation leads to a set of internal consistency conditions that link observed outcomes and productivity beliefs to the sequence of unobserved expected paths. Solving these conditions allows us to recover the full sequence of expected paths. The key intuition is that by focusing on how expected paths evolve in response to belief updates, our method effectively 'differences out' the influence of all time-varying deterministic fundamentals. This solution therefore disentangles the role of agents' beliefs from that of structural parameters in shaping observed outcomes.

Jointly solving the internal consistency conditions for all expected paths creates a prohibitively large fixed-point problem. We circumvent this complexity by recovering expected paths sequentially backward, transforming the high-dimensional problem into a single backward recursion. Each step of this recursion calculates the expected path from period t as a

[1] In developing the methodology, we assume researchers collect agents' beliefs on exogenous fundamental productivity over time; in one of our applications, we demonstrate how to leverage beliefs on endogenous outcomes when beliefs on fundamental productivity are not readily available.

deviation from the one from $t+1$. Crucially, solving for this deviation requires only the differences in beliefs between t and $t+1$; it does not require knowledge of the levels of underlying fundamentals like migration and trade costs. Furthermore, once the full sequence of expected paths is recovered, we can compute counterfactuals under alternative fundamental paths in deviations from any recovered path. Our approach therefore tractably generalizes the exact hat algebra method (Dekle et al., 2007; Caliendo et al., 2019) to dynamic settings with imperfect foresight.[2]

Section 2.6 addresses the second challenge: incorporating the dynamic impacts of aggregate uncertainty, which can arise from many locations and sectors. We develop a method to compute the second-order accurate solution based on the simulations of first-order solutions. The second-order approximation of the equilibrium yields a system of equations with numerous second-order terms (e.g., conditional variances and covariances of endogenous outcomes). Our key insight is that these terms can themselves be computed with second-order accuracy using only first-order simulations. This approach readily accommodates the flexible, time-varying, and location-specific stochastic processes (and associated beliefs) that are central to our framework, with no additional computational burden relative to standard, more restrictive processes (e.g., a common AR(1) process across locations).

Beyond these core challenges, environments without perfect foresight raise the possibility of heterogeneous beliefs across groups of agents—e.g., many people are still skeptical about climate change. Solving such models involves a well-known curse of dimensionality in higher-order beliefs (Townsend, 1983). In an extension of the model, we propose alternative structures of higher-order beliefs that ensure tractability while preserving the model's flexibility regarding fundamental processes and agents' beliefs. For clarity, the main text analyzes the two core challenges separately; the appendix demonstrates that the methods can be tractably combined to, for example, obtain counterfactuals under uncertainty given the observed data, or to accommodate these tractable forms of heterogeneous beliefs.

Finally, in applications of dynamic spatial models, a key interest lies in measuring the welfare effects of fundamental shocks. In settings with imperfect foresight, welfare measured ex-ante can differ from that measured ex-post (i.e., conditional on the realized path of fundamentals). We discuss in Section 2.7 how to recover both welfare measures from the

[2]Given the information on agents' beliefs about future productivity, an alternative approach is to simulate the model to search for the migration costs that rationalize the observed outcome series. This is computationally intractable because for any given level of migration costs, simulating the evolution of the economy under imperfect foresight requires solving multiple transition paths (as opposed to only one path under perfect foresight). The space for migration costs itself can also be large, especially if these costs vary across destination/origin and over time flexibly—as is often the case in applications of dynamic spatial models.

observed and counterfactual outcomes without needing to back out fundamentals.

In Sections 3 and 4, we apply our methodology to two empirical applications, both set in a geography comprising 50 U.S. states, 28 other countries, and multiple sectors. In the first application, we study the effects of China's productivity growth over 2001–2008 on the U.S. economy. We estimate a process characterizing the catch-up of China's manufacturing productivity with that of the U.S. We assume that agents make decisions based on their beliefs about the catch-up process and discipline these beliefs using manufacturing employment projections made by the Bureau of Labor Statistics, which started out overly optimistic but were revised downward gradually as the China shock unfolded. Interpreting these projections as reflecting the prevailing beliefs at the time, we infer that agents initially underestimated China's productivity catch-up but gradually, their beliefs converged to the realized path. We find that China's rapid productivity catch-up led to a smaller cumulative decline in U.S. manufacturing employment by 2008 and slightly larger aggregate welfare gains in the model with evolving beliefs compared to a counterfactual assuming agents had perfect foresight throughout the period. Behind the different welfare assessments are large differences in inferred migration costs—if we were to rationalize the same migration patterns under perfect foresight, the implicit cost of moving out of manufacturing would need to be about 50% higher.

Our second application is an ex-ante evaluation of how rising temperatures from 2014 to 2100 could affect welfare and the spatial allocation of economic activity in the United States. We focus on two salient departures from perfect foresight: first, agents face uncertainty regarding the path of future temperatures, and second, a substantial share of the population is skeptical about climate change. Compared to a baseline scenario with perfectly anticipated climate change, incorporating uncertainty leads to a moderately faster spatial reallocation from warmer states in the South to colder states in the North. The intuition stems from the convex nature of temperature damages: in already warmer locations (primarily the South), the marginal negative impact of further warming is greater. Consequently, a given amount of temperature uncertainty translates into a larger variance of future productivity in the South, accelerating out-migration towards cooler states. Incorporating uncertainty almost doubles the aggregate welfare losses for the U.S.—despite it being a second-order effect. The reason is that while the effects of perfectly anticipated warming can be positive for some states and negative for others—partially canceling out in aggregate—the additional welfare effects from uncertainty are negative everywhere. These losses accumulate to an aggregate impact comparable to that of perfectly anticipated climate change. Independent of aggregate uncer-

tainty, we incorporate climate skeptics (whose initial state-level distribution is informed by survey data). The presence of climate skeptics slows down spatial reallocation, resulting in an increase in the welfare of climate-change believers in the North and a decrease in their welfare in the South relative to the baseline scenario.

Our work contributes to and builds on several strands of the literature. First, it contributes to the growing body of research on spatial dynamics examining how agents' forward-looking decisions shape the economy's adjustment to shocks. This literature, reviewed in Desmet and Parro (2025), addresses topics such as the persistent effects of trade shocks (e.g., Artuç et al., 2010; Rodríguez-Clare et al., 2025; Dix-Carneiro et al., 2023), the impact of environmental changes (e.g., Balboni, 2025; Rudik et al., 2022; Cruz, 2023), and regional dynamics involving migration and investment (e.g., Kleinman et al., 2023; Giannone et al., 2023). Like many frameworks in this literature, our model is designed to bring spatial models with large state spaces to the data. Our contribution to this literature is twofold. Theoretically, we incorporate tractable departures from perfect foresight and demonstrate how to conduct counterfactuals and measure welfare under these conditions. Empirically, we apply this framework to quantify the role of belief dynamics and the impact of uncertainty and heterogeneous beliefs.

Second, our empirical findings add to two strands of the literature. Compared to existing work on the China shock (e.g., Autor et al., 2013; Caliendo et al., 2019; Galle et al., 2023; Rodríguez-Clare et al., 2025; Dix-Carneiro et al., 2023), we highlight the role of imperfect beliefs in shaping observed adjustment responses. Our second application complements recent work analyzing the effects of climate change in dynamic spatial models or integrated assessment models (e.g., Cruz, 2023; Cruz and Rossi-Hansberg, 2024; Rudik et al., 2022). We contribute by using a dynamic model with spatial granularity to quantify the impact of uncertainty and heterogeneous beliefs regarding future temperature rise.

Third, our approximation approach is connected to Kleinman et al. (2023), who conduct first-order approximations around steady states to derive analytical insights into the transition path in dynamic spatial models. Related are contemporary studies that incorporate aggregate uncertainty into spatial models using alternative approaches such as global approximation based on neural networks (Pang and Pin, 2022) or splines (Egger et al., 2024), and local perturbation around the steady state in continuous time (Bilal, 2023).[3]

[3]Both the global and local approaches have precedents in macroeconomics. See, e.g., Krusell and Smith (1998) for global solutions of models with aggregate uncertainty, and Schmitt-Grohé and Uribe (2004) and Kim et al. (2008) for the second-order perturbation method in solving DSGE models. The challenge in spatial equilibrium models is that state spaces are often larger and there can be many prices to clear many markets.

Also, Porcher (2022) studies a setting where productivity follows an AR(1) process and agents, facing information frictions on regional outcomes, solve a signal extraction problem with Gaussian noise. Instead of approximating the state-space solution of the model as in all these works, we approximate the model around deterministic *paths* and develop a simulation-based method for the second-order solutions. This enables us to accommodate generic stochastic processes across locations and flexible, history-dependent belief formation that are central to many applications in dynamic spatial economies.[4] More importantly, beyond incorporating uncertainty, another main focus of our paper is to develop a tractable approach to recover counterfactual outcomes and measure welfare when the observed data are generated by agents without perfect foresight.

Fourth, our use of external data to recover agents' beliefs connects it to studies exploring alternative ways to infer evolving beliefs, such as inferring them from observed decisions (e.g., Dickstein and Morales, 2018; Bombardini et al., 2023; Porcher et al., 2025). More broadly, it also relates to the literature in macro and international economics emphasizing the role of learning or expectations regarding fundamentals such as productivity (Kozlowski et al., 2020; Bui et al., 2022) and trade barriers (e.g., Alessandria et al., 2025).

2 A Dynamic Stochastic Spatial Model

2.1 Economic Environment

The economy consists of N locations, denoted by n or i. Time is discrete and indexed by $t = 1, ..., T$, where T can be finite or infinite. In each location and period, a continuum of individuals work, consume, and make forward-looking migration decisions subject to migration frictions and idiosyncratic location-specific taste shocks. Locations trade with each other. To simplify notation, we present a one-sector model here; our results generalize to a multi-sector extension with intermediate inputs, which we use in applications.

Each location n is endowed with initial labor l_{n1} and characterized by a set of economic fundamentals $\{z_{nt}, m_{nit}, \kappa_{nit}\}$, where z_{nt} is fundamental productivity, and m_{nit} and κ_{nit} are bilateral migration and trade frictions, respectively. These fundamentals can be time-varying. Without loss of generality, we assume that the evolution of trade and migration

[4] Focusing on a sequence-space model, Auclert et al. (2021) shows how to recover numerical gradients over the transition path. In our model, we show that the sequence-space gradients and Hessians can be expressed in closed-form as a function of equilibrium objects. Given these objects, our main contribution lies in leveraging them to address the challenges of imperfect foresight and uncertainty in ex-ante and ex-post counterfactuals.

costs is deterministic and perfectly anticipated, while the evolution of productivity z_{nt} is stochastic. Our methodology extends to settings in which m_{nit} or κ_{nit} are also stochastic.

2.2 Stochastic Fundamentals and Beliefs

We now describe the evolution of z_{nt}. Let $z_t \equiv (z_{1t}, z_{2t}, ..., z_{Nt})$ be the vector of productivity across locations in period t, and let $z^t \equiv \{z_s\}_{s=1}^t$ denote the history of these vectors up to period t. Let Θ denote the set of possible outcomes for z_t; then, the set of possible histories for z^t up to t is Θ^t. For example, if $z_{nt} \in R_+$, then $\Theta = R_+^N$ and $\Theta^t = R_+^{N\times t}$.

Productivity processes. The vector of fundamental productivity z_t evolves according to a probability density function (pdf) $g_t(z_{t+1}|z^t)$, which captures how the history up to t determines the likelihood of outcomes in $t+1$. This general notation places no restrictions on whether productivity in different locations evolves independently or not; with an appropriate choice of g_t, it can accommodate arbitrary (and potentially non-stationary, time-varying) dependence of the future productivity path on past productivity.[5] Such generality is useful in many applications of dynamic spatial models, as we show in our two applications.

Once the g_t function is specified, the likelihood of all future productivity paths is determined. In particular, given the history z^t, the likelihood of a path $z^{t'} \in \Theta^{t'}$ for $t' > t$ is expressed as $g(z^{t'}|z^t) \equiv \prod_{i=0}^{t'-t-1} g(z_{t+i+1}|z^{t+i})$. For simplicity, we have suppressed the subscript t in g_t. This is without loss of generality given that the dependence of $g_t(z_{t+1}|z^t)$ on t can be subsumed in z^t.

Agent's belief. Agents make forward-looking decisions based on their beliefs about future productivity. We denote by $f(z_{t+1}|z^t)$, short for $f_t(z_{t+1}|z^t)$, the conditional pdf of the vector z_{t+1} from the perspective of agents in period t given a history z^t. Beliefs are common among agents. We later extend our methodology to incorporate beliefs that differ among groups of agents. Following the notation earlier, the likelihood of a productivity path according to belief f is $f(z^{t'}|z^t) \equiv \prod_{i=0}^{t'-t-1} f(z_{t+i+1}|z^{t+i})$.

By choosing f suitably, our framework nests a broad range of belief formulations. The special case where $f = g$ corresponds to rational expectations. When $f \neq g$, it may be due to agents' not knowing the functional form of g or their gradual learning about its

[5] For example, the setup accommodates processes of the following form: $z_{t+1} = a_t + \sum_{t'=0}^{t} b_{t'} z_{t'} + \epsilon_{t+1}$, with $\epsilon_{t+1} \sim N(0, \sigma_{t+1}^2)$ being the innovation term, and a_t, $\{b_{t'}\}_{t'=0}^t$, and σ_{t+1} being time-varying coefficients. These time-varying coefficients can themselves be a function of z^t. As will become clear later, longer and more flexible history dependence does not increase the computational burden in our approach.

parameters.[6] Different forms of learning can be accommodated by specifying how f varies with new information. For example, learning can be Bayesian or myopic (extrapolating); regardless of the case, agents could be 'naive' learners as defined by the anticipated utility framework developed by Kreps (1998), thinking that their belief of the world is correct and will not change, only to discover in the next period that the belief is incorrect and then update it. Alternatively, they could be 'sophisticated' in the sense that they anticipate revising their beliefs as future data arrive. The latter case typically poses computational challenges because agents' beliefs about their own future belief revisions become part of the state space (Cogley and Sargent, 2008), but it can be accommodated in our setup.[7]

Like the true g process, the belief process f can be time-varying and history dependent. In settings where agents need to form beliefs about time-varying and history dependent g, such flexibility is essential. Even when the g processes are simple (e.g., an AR(1) process), this flexibility is valuable because it accommodates realistic features of beliefs. For instance, agents may be more uncertain about fundamentals in the more distant future (e.g., future temperature rise); they may engage in history/context-dependent learning, in which case their reactions to new information depend on history. Such mechanisms have gained empirical support (e.g., Greenwood and Shleifer, 2014; Malmendier and Nagel, 2016).

Finally, in our notation above, agents form beliefs based on exogenous past fundamentals. If agents instead form beliefs about future fundamentals based on endogenous past outcomes that depend on the realized fundamentals, then their beliefs are a composite function in the realized fundamentals. For example, $\tilde{f}\left(z_{t+1}|GDP^t(z^t)\right)$, where GDP^t is a function of z^t. Such beliefs can be expressed in our notation by defining $f(z_{t+1}|z^t) \equiv \tilde{f}\left(z_{t+1}|GDP^t(z^t)\right)$.[8,9]

[6] For example, let the true process be $g(z_{t+1}|z^t,\phi)$, where ϕ represents some parameters in g. The agents' perceived process might be $f(z_{t+1}|z^t,\hat{\phi})$, differing either because of $f \neq g$ or because $\hat{\phi} \neq \phi$.

[7] Consider an environment described in Footnote 6. In our setup, the likelihood of a future path $z^{t'}$ perceived in period 0 by naive agents is $f(z^{t'}|z^t,\hat{\phi}) = \prod_{i=0}^{t'-t-1} f(z_{t+i+1}|z^{t+i},\hat{\phi}(z^0))$; for sophisticated agents, this likelihood is instead $f(z^{t'}|z^t,\hat{\phi}) = \prod_{i=0}^{t'-t-1} f(z_{t+i+1}|z^{t+i},\hat{\phi}(z^{t+i}))$, where agents in period 0 anticipate updating $\hat{\phi}$ with new data.

[8] This also nests the case when GDP is a noisy function of realized productivity—due to, for example, measurement errors in statistics, in which case agents' forecast involves solving a signal extraction problem.

[9] Another alternative is to specify beliefs about *endogenous* outcomes. For example, agents might form beliefs as $f(GDP_{t+1}|z^t)$. The drawback of this approach is that it does not easily lead to general equilibrium counterfactuals for hypothetical fundamental processes. For instance, to study how U.S. manufacturing employment would have evolved in the absence of the China shock, one would need to specify how beliefs on local GDP evolve in a scenario with slower Chinese productivity growth. In contrast, under our specification, once $f(z_{t+1}|z^t)$ is disciplined using the data, beliefs in counterfactual scenarios are readily defined.

2.3 Forward-Looking Decisions with Stochastic Fundamentals

The continuum of agents in the economy consume in their current location and make forward-looking migration decisions subject to migration frictions m_{nit} and idiosyncratic taste shocks ϵ_{it}. We assume that the taste shocks ϵ_{it} are orthogonal to the fundamental productivities and drawn *i.i.d* from a Type I Extreme Value (Gumbel) distribution with dispersion parameter ν.

Given the initial worker distribution l_{n1}, the state of the economy in period t is summarized by the productivity history z^t. Conditional on z^t, the value of a location n for an agent with a realization of idiosyncratic taste shocks $\epsilon_t \equiv (\epsilon_{1t}, \epsilon_{2t}, ..., \epsilon_{Nt})$ is given by

$$V_{nt}(z^t, \epsilon_t) = U(c_{nt}(z^t)) + \max_{\{i\}_{i=1}^N} \{\beta \mathbb{E}[V_{it+1}(z^{t+1}, \epsilon_{t+1})|z^t] - m_{nit} + \nu \epsilon_{it}\}, \tag{1}$$

where $U(c_{nt}(z^t))$ is the flow utility of the agent at location n, and $c_{nt}(z^t)$ is the real income defined as the ratio between wage $w_{nt}(z^t)$ and price index $P_{nt}(z^t)$. The conditional expectation in $\mathbb{E}[V_{it+1}(z^{t+1}, \epsilon_{it+1})|z^t]$ is taken over both the future taste shocks ϵ_{t+1} and future productivity z_{t+1}. Formally,

$$\underbrace{\mathbb{E}[V_{it+1}(z^{t+1}, \epsilon_{it+1})|z^t]}_{\equiv v_{it+1}(z^t)} = \int_\Theta \underbrace{[\int_{\epsilon_{t+1}} V_{it+1}(z^{t+1}, \epsilon_{t+1}) d\mathcal{H}(\epsilon_{t+1})]}_{\equiv v_{it+1}(z^{t+1})} f(z_{t+1}|z^t) dz_{t+1}, \tag{2}$$

where $\mathcal{H}(\epsilon_{t+1})$ is the cumulative distribution function of $\epsilon_{t+1} \equiv (\epsilon_{1t+1}, \epsilon_{2t+1}, ..., \epsilon_{Nt+1})$.

The notation in equation (2), $v_{it+1}(z^t)$ and $v_{it+1}(z^{t+1})$, distinguishes between two different conditional expectations. In particular, $v_{it+1}(z^{t+1})$ is the expected value conditional on a history z^{t+1} (integrating only over ϵ_{t+1}), whereas $v_{it+1}(z^t)$ is the expected value from the perspective of period t (integrating over both ϵ_{t+1} and future fundamental shocks z^{t+1}).

Integrating equation (1) over the idiosyncratic taste shocks ϵ_t and applying the properties of the Gumbel distribution yields $v_{nt}(z^t)$, which is the average location value of agents in n at time t after a history z^t, and the share of these agents migrating to i:

$$v_{nt}(z^t) \equiv \int_{\epsilon_t} V_{nt}(z^t, \epsilon_t) d\mathcal{H}(\epsilon_t) = U(c_{nt}(z^t)) + \nu \ln(\sum_{i=1}^N \exp(\beta v_{it+1}(z^t) - m_{nit})^{1/\nu}), \tag{3}$$

$$\mu_{nit}(z^t) = \frac{\exp\left(\beta v_{it+1}(z^t) - m_{nit}\right)^{1/\nu}}{\sum_{h=1}^N \exp(\beta v_{ht+1}(z^t) - m_{nht})^{1/\nu}}. \tag{4}$$

Migration decisions at time t are thus shaped by the expected continuation value $v_{it+1}(z^t)$.[10]

[10] See Online Supplement C.1 for the full derivation of equations (3) and (4).

Migration shares, together with the distribution of workers at the beginning of period t, determined in $t-1$ and denoted by $l_{it}(z^{t-1})$, govern the law of motion for labor:

$$l_{nt+1}(z^t) = \sum_{i=1}^{N} \mu_{int}(z^t) l_{it}(z^{t-1}). \tag{5}$$

We model the labor demand with a general equilibrium trade structure with CES preferences following Eaton and Kortum (2002). The assumed structure implies that the share of goods that location n purchases from location i, denoted by $\lambda_{nit}(z^t)$, is

$$\lambda_{nit}(z^t) = \frac{z_{it}[w_{it}(z^t)\kappa_{nit}]^{-\theta}}{\sum_{h=1}^{N} z_{ht}[w_{ht}(z^t)\kappa_{nht}]^{-\theta}} \tag{6}$$

where θ is the trade elasticity, κ_{nit} is the trade cost, and $w_{it}(z^t)$ is the wage. The consumer price index and the labor market clearing condition are given by, respectively,

$$P_{nt}(z^t) = \Gamma\Big[\sum_{i=1}^{N} z_{it}[w_{it}(z^t)\kappa_{nit}]^{-\theta}\Big]^{-1/\theta}, \tag{7}$$

$$\text{and} \quad w_{nt}(z^t) l_{nt}(z^{t-1}) = \sum_{i=1}^{N} \lambda_{int}(z^t) w_{it}(z^t) l_{it}(z^{t-1}), \tag{8}$$

where Γ is a constant. Given a labor allocation, equations (6) and (8) characterize the static trade equilibrium. We can now define the stochastic sequential equilibrium of the economy.

Definition 1. ***A stochastic sequential equilibrium of the model*** *is a set of state-contingent prices* $\{w_{nt}(z^t), P_{nt}(z^t)\}_{n=1,t=1}^{N,T}$, *location values* $\{v_{nt}(z^t)\}_{n=1,t=1}^{N,T}$, *and allocations of goods and labor* $\{\lambda_{nit}(z^t), \mu_{nit}(z^t), l_{nt}(z^{t-1})\}_{n=1,i=1,t=1}^{N,N,T}$ *that satisfies the value function (3), migration shares (4), the law of motion for labor (5), the trade shares (6), the price index (7), and the labor market clearing condition (8).*

Two remarks are in order. First, unlike in standard models where the state is summarized by current productivity z_t and labor allocation $\{l_{it}\}$, in our framework the state variable is the entire productivity history z^t. This dependence on history arises because not only can $\{\kappa_{nit}, m_{nit}\}$ vary arbitrarily over time, but more importantly, the belief and true productivity processes (f and g) can have history dependence. This feature enables the model to accommodate general productivity and belief-updating processes.

Second, the presence of forward-looking decisions and history dependence implies that to obtain the decision rule in any period t, one must solve the full set of equations for *all* periods $t' > t$ and all possible future trajectories $z^{t'}$, which is infeasible. We therefore

implement a local approximation method. Building on this approximation, we address the core challenges of this paper: first, we describe how to recover counterfactual outcomes when agents do not have perfect foresight; second, we explain how to incorporate aggregate uncertainty in spatial settings with flexible belief and productivity processes; and third, we show how this approach can be extended to incorporate heterogeneous beliefs.

2.4 Local Approximation of the Stochastic Equilibrium

To solve the model, we conduct a local approximation of the stochastic equilibrium around a deterministic sequential equilibrium path. We take this deterministic path as given for now and explain later how to construct the path for different exercises.

Suppose the economy is at period t after history z^t. Solving for the migration decision at t requires knowing the outcomes in all possible future states. Up to the first order, however, the decision depends only on the expected future path of the economy (in deviations from the deterministic path). To formalize this, let $x_{t'}(z^{t'})$ be a generic outcome x (e.g., wages, location values) at a future time t'; $\bar{x}_{t'}$ be the corresponding outcome along the deterministic path; and $\hat{x}_{t'}(z^{t'})$ be the deviation of $x_{t'}(z^{t'})$ from $\bar{x}_{t'}$. These deviations represent level differences for location values (v_{nt}) and log differences for other variables.

For any variable $x_{t'}(z^{t'})$ in any period $t' = t, t+1, \ldots, T$, we take a first-order Taylor expansion of the equilibrium conditions (equations (3) to (8)) around $\bar{x}_{t'}$ to obtain $\hat{x}_{t'}(z^{t'})$. Integrating $\hat{x}_{t'}(z^{t'})$ using the likelihood $f(z^{t'}|z^t)$ delivers the following system of linear equations in expected deviations.

$$\text{Location value:} \quad \mathbb{E}_t \hat{v}_{nt'} = \frac{\partial U(\bar{c}_{nt'})}{\partial \ln \bar{c}_{nt'}} \mathbb{E}_t \hat{c}_{nt'} + \beta \sum_{i=1}^{N} \bar{\mu}_{nit'} \mathbb{E}_t \hat{v}_{it'+1} \tag{9}$$

$$\text{Migration flows:} \quad \mathbb{E}_t \hat{\mu}_{nit'} = \beta/\nu \sum_{m=1}^{N} \big(\mathbb{1}(m=i) - \bar{\mu}_{nmt'}\big) \mathbb{E}_t \hat{v}_{mt'+1} \tag{10}$$

$$\text{Labor distribution:} \quad \mathbb{E}_t \hat{l}_{nt'+1} = \sum_{i=1}^{N} \frac{\bar{\mu}_{int'} \bar{l}_{it'}}{\bar{l}_{nt'+1}} \big(\mathbb{E}_t \hat{\mu}_{int'} + \mathbb{E}_t \hat{l}_{it'}\big) \tag{11}$$

$$\text{Trade shares:} \quad \mathbb{E}_t \hat{\lambda}_{nit'} = -\theta(\mathbb{E}_t \hat{w}_{it'} - \mathbb{E}_t \hat{P}_{nt'}) + \mathbb{E}_t \hat{z}_{it'} \tag{12}$$

$$\text{Price index:} \quad \mathbb{E}_t \hat{P}_{nt'} = \sum_{i=1}^{N} \bar{\lambda}_{nit'} \big(\mathbb{E}_t \hat{w}_{it'} - 1/\theta \cdot \mathbb{E}_t \hat{z}_{it'}\big) \tag{13}$$

$$\text{Market clearing:} \quad \mathbb{E}_t \hat{w}_{nt'} + \mathbb{E}_t \hat{l}_{nt'} = \sum_{i=1}^{N} \frac{\bar{\lambda}_{int'} \bar{w}_{it'} \bar{l}_{it'}}{\bar{w}_{nt'} \bar{l}_{nt'}} \big(\mathbb{E}_t \hat{\lambda}_{int'} + \mathbb{E}_t \hat{w}_{it'} + \mathbb{E}_t \hat{l}_{it'}\big). \tag{14}$$

Here, $\mathbb{E}_t[\cdot]$ denotes the expectation at t for a future variable in t'. For example, $\mathbb{E}_t\hat{v}_{nt'} \equiv \int_{\Theta^{t'}} \hat{v}_{nt'}(z^{t'}) f(z^{t'}|z^t) dz^{t'}$. This system of equations holds for $t' = t, t+1, \ldots, T$. It takes two inputs. The first one is the sequence of expected productivity deviations from the deterministic path $\{\mathbb{E}_t\hat{z}_{t'}\}_{t'=t}^{T}$, where $\mathbb{E}_t\hat{z}_{t'} \equiv \int_{\Theta^{t'}} [\ln(z_{t'}) - \ln(\bar{z}_{t'})] f(z^{t'}|z^t) dz^{t'}$. Here, by constructing $f(z^{t'}|z^t)$ appropriately, these expected deviations accommodate different models of belief formation such as rational expectations and naive or sophisticated learning. The term $\{\mathbb{E}_t\hat{z}_{t'}\}_{t'=t}^{T}$ is zero under rational expectations but more generally, it depends on the history up to t. The exact belief formation process, encoded by the function f, can be disciplined by data, as described in the next section and illustrated in our first application. The second input is the deterministic equilibrium path, $\{\bar{x}_{t'}\}_{t'=t}^{T} \equiv \{\bar{c}_{it'}, \bar{w}_{it'}, \bar{l}_{it'}, \bar{\mu}_{int'}, \bar{\lambda}_{nit'}\}_{i=1,n=1,t'=t}^{N,N,T}$. These outcomes reflect the influence of fundamentals, such as trade and migration costs.

Solving this system of equations using standard algorithms (see Appendix A.1 for an example) delivers the expected path of all future endogenous outcomes from the perspective of agents at t, i.e., $\left\{\mathbb{E}_t\hat{v}_{it'}, \mathbb{E}_t\hat{\mu}_{nit'}, \mathbb{E}_t\hat{l}_{it'}, \mathbb{E}_t\hat{w}_{it'}, \mathbb{E}_t\hat{\lambda}_{nit'}, \mathbb{E}_t\hat{P}_{it'}\right\}_{i=1,n=1,t'=t}^{N,N,T}$. The first-period outcomes in the expected path ($t' = t$) constitute a first-order approximation to the actual outcomes in t given history z^t. In particular, the actual migration choices by agents in t are obtained as: $\mu_{nit}(z^t) \approx \bar{\mu}_{nit} \cdot \exp(\mathbb{E}_t\hat{\mu}_{nit})$. Therefore, solving these equations recovers agents' migration choices in t, which determine the labor allocation $l_{it+1}(z^t)$ for $t+1$. We summarize these aspects of our approximate solution in the following proposition:

Proposition 1. *Given the labor allocation in a period $t \in 1, .., T$; the deterministic path, $\{\bar{x}_{t'}\}_{t'=t}^{T}$; and the expected productivity deviations $\{\mathbb{E}_t\hat{z}_{t'}\}_{t'=t}^{T}$, the system of linear equations (9) to (14) determines the actual choices in t, $\{x_t(z^t)\}$, as well as the expected path of endogenous variables (in deviations), $\{\mathbb{E}_t\hat{x}_{t'}\}_{t'=t}^{T}$.*

Proof. See Appendix A.1. □

This approximation provides a mapping between productivity beliefs, the expected path of outcomes, and the actual evolution of the economy. We build on this local approximation for several purposes. One such purpose is to simulate the evolution of an economy for given paths of future fundamentals. Starting from an initial labor allocation, solving the system delivers the migration shares $\mu_{nit}(z^t)$, which determine the labor allocation for $t+1$. Repeating this process recursively for $t+1, t+2, \ldots, T$, traces out the economy's evolution under a given fundamental path z^T and belief process f. We refer to such exercises, where the evolution of the economy has not yet realized, as 'ex-ante counterfactuals.'

Instead of simulating forward an economy under assumed paths of fundamentals and beliefs, in the next subsection, we focus on a different and more challenging task: performing ex-post counterfactuals. In such exercises, a researcher observes the realized evolution of the economy and aims to determine the alternative evolution that would have occurred under a particular shock. We explain the identification challenge in conducting ex-post counterfactuals when agents lack perfect foresight and propose a tractable solution.

From the preceding discussion, provided that the deviation terms are defined consistently, any deterministic path can serve as the approximation path. In practice, choosing a path that is 'close' to agents' expectations reduces approximation error. Consider the limit case where the deterministic path starts from the same initial labor allocation as the stochastic economy and has a productivity sequence that coincides with agents' beliefs, i.e., $\{l_{it}(z^t)\} = \{\bar{l}_{it}\}$ and $\{\bar{z}_{t'}\}_{t'=t}^{T} = \{\mathbb{E}_t z_{t'}\}_{t'=t}^{T}$. In this case, the deterministic path is itself a first-order approximation to the expected path—a familiar certainty equivalence result.

This result implies that for ex-ante counterfactuals, one could solve the deterministic nonlinear model directly to obtain a first-order solution. However, our more general approximation approach is still essential for two reasons. First, as we show below, it can be extended to obtain second-order accurate solutions that account for uncertainty around the path of future fundamentals. Second, this second-order solution adapts naturally to ex-post counterfactuals, in which the observed outcomes are shaped by expected but unrealized future fundamentals.

Comparison to steady-state-based approximation. Although the approximation of the system itself is not the main outcome in our paper, it is worthwhile comparing it with alternative local solution methods. A standard local solution method in macroeconomics (Uhlig, 1995) involves approximating around the steady state. Approximating around a deterministic path has two useful properties for conducting counterfactuals in spatial economies.

First, it accounts for the fact that in many applications of dynamic spatial models, the economy may begin from a point with a drastically different economic landscape from both the pre-shock and post-shock steady states. This could be due to the size of the shock under investigation being large or other forces not directly related to the shock (e.g., the initial distribution of agents being far from the steady-state distribution). Using either steady state to approximate agents' decisions can lead to substantial approximation errors. One might think that such approximation errors are of second order, but it is important to be precise about the sense in which they are 'second order.'

For concreteness, consider the approximation of a smooth *single-variable function* $f(x)$.

Suppose we are interested in finding $f(x_2+\delta)-f(x_2)$, i.e., the impact of a change of δ on $f(x)$ around x_2. A first-order approximation around x_2 gives us $f(x_2+\delta)-f(x_2) = f'(x_2)\cdot\delta+O(\delta^2)$, with the error being second order in δ, the shock of interest. If we approximate around x_1 instead, we obtain

$$\begin{aligned} f(x_2+\delta)-f(x_2) &= f'(x_2)\cdot\delta+O(\delta^2) = f'(x_1)\cdot\delta+[f'(x_2)-f'(x_1)]\cdot\delta+O(\delta^2) \\ &= f'(x_1)\cdot\delta+f''(x_1)\cdot(x_2-x_1)\cdot\delta+O\left((x_2-x_1)^2\cdot\delta\right)+O(\delta^2). \end{aligned}$$

Therefore, approximating around x_1 introduces an additional error term, $f''(x_1)\cdot(x_2-x_1)\cdot\delta+O((x_2-x_1)^2\cdot\delta)$. Following this analogy, if the steady state (x_1) is far away from the actual path (x_2), this error is first order in the shock δ.[11] Approximating around a suitably chosen path reduces such errors.[12]

Second, our approximation is based on a within-period, cross-economy comparison between the stochastic economy and a deterministic one. By choosing the approximation path appropriately, time-varying fundamentals such as trade and migration costs are accounted for directly. While this feature may not provide much utility for ex-ante counterfactuals, where the researcher is free to take a stance on the evolution of these fundamentals, its value becomes apparent in ex-post counterfactuals. As we show below, when recovering counterfactuals from observed data, the researcher must take these time-varying fundamentals as given, since they have already shaped the observed outcomes.

2.5 Ex-Post Counterfactual without Perfect Foresight

The local approximation described in the previous section can be used for ex-ante evaluations of hypothetical *future* changes in productivity (or beliefs), taking as given an assumed path for migration and trade costs. Equally relevant, however, is studying the effects of *past* changes in fundamentals. Suppose, after observing a time series of outcomes, a researcher seeks to recover a counterfactual evolution of the economy. This allows one to ask important

[11] The error term $(x_2-x_1)\cdot\delta$ is second-order in δ if x_2 itself deviates from x_1 only because of shocks of order δ. This may be a reasonable assumption in business cycle analysis, where all deviations from the steady state are driven by such shocks. In dynamic spatial applications, however, the economy can be far from its steady state due to other time-varying fundamentals or its initial conditions. In such cases, the distance x_2-x_1 can be much larger than δ, making the approximation error significant.

[12] Consistent with this view, Mennuni et al. (2024) show that perturbation around a deterministic path is numerically more accurate than around the steady state. We differ from Mennuni et al. (2024) in two important aspects. First, instead of perturbation, we use local approximations in the sequence space. We discuss in Section 2.6 the advantages of the two approaches. Second, we build on path approximation for ex-post counterfactuals and for incorporating uncertainty and heterogeneous beliefs.

questions such as: What is the labor market effect of a trade liberalization episode? What are the welfare effects of the China shock? What is the role of a transport infrastructure project (which may or may not materialize as planned) in shaping a city's development?

In these types of counterfactuals, observed data are taken as given, and the model fundamentals shaping the data, such as migration costs, must be disciplined by the data internally—either implicitly through model inversion (e.g., Caliendo et al., 2019) or explicitly through structural estimation (e.g., Dix-Carneiro, 2014). Regardless of the approach, the assumptions on agents' beliefs are crucial. For example, if few American workers moved out of manufacturing in response to the persistent increase in import competition from China, it could be because they did not anticipate the increase to be so large and persistent or because they faced large migration costs. Whether it is the former, the latter, or a combination of both factors, will shape the counterfactual outcomes.

Carrying out such counterfactuals therefore requires disentangling the role of beliefs about future productivity from that of other fundamentals in shaping agents' sequential decisions. Two challenges arise. First, disciplining deviations from perfect foresight requires data on beliefs. This challenge is alleviated by recent progress in collecting data on beliefs about fundamentals or outcomes, such as surveys measuring the expectations of households, businesses, and financial analysts (e.g., Armantier et al., 2017; Coibion et al., 2018) or newspaper coverage (Bui et al., 2022).[13] More generally, even if belief data on the full set of fundamentals or outcomes across all geographic units and sectors are unavailable, it is often possible to gather beliefs on productivity in a subset of sectors or regions or on endogenous outcomes, such as expectations of GDP or employment, which contain information about productivity beliefs. We illustrate how to extract such information in the first application.

Armed with belief data, the second challenge is to disentangle the role of beliefs from other fundamentals in a theory-consistent way. To illustrate, consider the economy depicted in Figure 1. Agents in the economy form beliefs about the productivity path and make migration decisions accordingly. Figure 1a shows the realized productivity path and agents' beliefs at different periods, which are overly optimistic. Figure 1b displays the expected evolution of an economic outcome (e.g., labor allocation) according to agents' beliefs alongside the actual evolution. A researcher observes the realized path of endogenous outcomes from $t = 1$ to T but not the sequence of expected paths. The researcher aims to examine how much of the change in the economic outcome is due to the increase in productivity, so they

[13] See also Porcher et al. (2025), who imposes structure on how migration costs vary over time and exploits the curvature in migration payoff functions to identify agents' beliefs based on migration decisions.

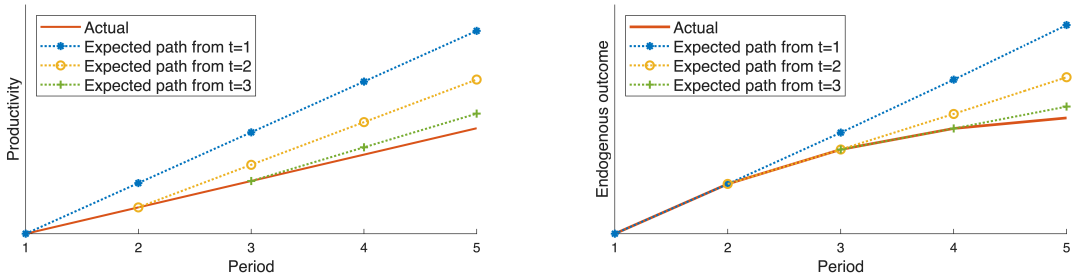

(a) Productivity (b) Endogenous Outcome

Figure 1: An Economy Under Evolving Beliefs

Note: Illustration of an economy in which a location receives positive productivity shocks and agents in the model are too optimistic about the shocks. Panel (a) presents the productivity path in the location receiving positive productivity shocks. The figure shows the actual productivity path (red solid line) and the beliefs about the productivity path formed in each period (dotted lines with markers). Panel (b) presents the actual path of an endogenous outcome (e.g., labor allocation) in the location and belief about that outcome from each period.

seek to construct a counterfactual evolution of the economy under constant productivity.

To construct this counterfactual, the researcher must discipline agents' migration costs, which influence their forward-looking migration decisions. In perfect foresight environments, the realized outcomes can be used to recover migration costs (Artuç et al., 2010, Caliendo et al., 2019, Dix-Carneiro et al., 2023). In environments without perfect foresight, however, decisions are shaped not by the realized outcomes, but by agents' expectation of these outcomes. Therefore, the role of beliefs must be accounted for. An intuitive strategy would be: (1) measure external beliefs on fundamentals; (2) for any given set of migration costs, simulate the economy's evolution; and (3) choose the migration costs that make the simulated evolution align with the observed data. While intuitive, this approach is intractable. Because agents' beliefs evolve over time, to simulate just one evolution (the solid line in Figure 1b), one must solve for T different transitional paths (the dashed lines). This entire process must then be repeated in a search over the space of migration costs to identify the ones consistent with actual evolution, a problem that becomes substantially more challenging if costs are time-varying or region-specific, as is often the case in applications of dynamic spatial models.[14]

A recursive solution. We develop a recursive solution to tackle this problem, leveraging two key insights. First, agents' expected paths of future endogenous outcomes (e.g., wages, allocations, as summarized by the dashed lines in Figure 1b) determine their forward-

[14]The complexity increases further if data on productivity beliefs are incomplete or unavailable. For example, if belief data cover only productivity in a subset of locations or sectors, or only endogenous outcomes such as output, then the procedure may need to be nested within an outer layer that parametrizes a belief function to match the available belief information, adding another layer of computational burden.

looking decisions. Therefore, these expected paths contain information about the expected path of fundamentals. Once these expected paths have been recovered, the information they reveal can be used to solve for counterfactual paths. The expected paths of outcomes are of course unobserved. The second insight we use is that, the difference between what agents in period t expect to do in $t+1$, $t+2$, ... (dashed lines in the figure) and what they actually do in these periods (solid line) arises entirely due to deviations between the expected and realized productivity. This creates a set of internal consistency conditions that link the sequence of unobserved expected paths to the observed evolution of beliefs and actual outcomes. Because aside from productivity beliefs, all expected paths are conditioned on the same underlying fundamentals (such as migration costs), their difference over time allows us to 'difference out' these fundamentals, isolating their role in shaping observed outcomes from the role of time-varying beliefs.

Solving these conditions simultaneously for all expected paths creates a prohibitively large fixed-point problem. We circumvent this by recovering the expected paths sequentially backward, transforming the high-dimensional problem into a single backward recursion. We now describe the intuition of this recursive solution, deferring the details to Appendix A.2.

Suppose for now that the entire sequence of productivity beliefs $\{\mathbb{E}_t z_{t'}\}_{t=1,t'\geq t}^{T}$ is observed. If the researcher has recovered the expected future outcomes according to the beliefs in period t, i.e., $\{\mathbb{E}_t \mu_{nit'}, \mathbb{E}_t l_{it'}, \mathbb{E}_t w_{it'}, \mathbb{E}_t \lambda_{nit'}, \mathbb{E}_t P_{it'}\}_{i=1,n=1,t'\geq t}^{N,N,T}$, which we term 'the expected path from t' and denote by $\{\mathbb{E}_t x_{t'}\}_{t'\geq t}$, then the expected path from $t-1$, $\{\mathbb{E}_{t-1} x_{t'}\}_{t'\geq t-1}$, can be derived in two steps.

Step 1: recover the expected path from $t-1$ for periods t onward. We first solve for $\{\mathbb{E}_{t-1} x_{t'}\}_{t'\geq t}$ by approximating around the known path $\{\mathbb{E}_t x_{t'}\}_{t'\geq t}$. We use the system of equations (9)–(14), setting the deterministic approximation path in these equations to the known path $\{\mathbb{E}_t x_{t'}\}_{t'\geq t}$ and the 'shock' to be the belief update between $t-1$ and t: $\{\mathbb{E}_{t-1} \ln z_{t'} - \mathbb{E}_t \ln z_{t'}\}_{t'\geq t}$. The solution to this system is the deviations between the two expected paths, which allows us to recover the level of the expected path from $t-1$ for periods t and onward: $\{\mathbb{E}_{t-1} x_{t'}\}_{t'\geq t}$

Step 2: recover the outcome at $t-1$. The solution from Step 1 delivers the expected path *from t onward*, but not at $t-1$ because the approximation path starts only at t. To complete the expected path, we use the fact that the actual migration and other outcomes at time $t-1$ must be consistent with expectations formed at $t-1$, that is, $\mathbb{E}_{t-1} x_{t-1} = x_{t-1}$. We can therefore append the observed outcome $\{x_{t-1}\}$ to the sequence recovered in Step 1 to obtain the full expected path from $t-1$, $\{\mathbb{E}_{t-1} x_{t'}\}_{t'\geq t-1}$.

These two steps recover $\{\mathbb{E}_{t-1}x_{t'}\}_{t'\geq t-1}$ from $\{\mathbb{E}_t x_{t'}\}_{t'\geq t}$. By iterating these steps backward from the final period T, we can recover the entire sequence of expected paths: $\{\mathbb{E}_{T-1}x_{t'}\}_{t'\geq T-1}$, $\{\mathbb{E}_{T-2}x_{t'}\}_{t'\geq T-2}$,..., until $\{\mathbb{E}_1 x_{t'}\}_{t'\geq 1}$. To start the recursion, note that the expected path from period T is simply the realized outcome in that period, $\{x_T\}$. We can therefore recover $\{\mathbb{E}_{T-1}x_{t'}\}_{t'=T}$ by approximating around $\{x_T\}$ with the final belief error, $\{\mathbb{E}_{T-1}\ln z_T - \ln z_T\}$, as the shock. Analogous to our previous discussion, appending $\{\mathbb{E}_{T-1}x_{t'}\}_{t'=T}$ with the actual outcome at $T-1$ yields $\{\mathbb{E}_{T-1}x_{t'}\}_{t'\geq T-1}$.

This recursive process achieves tractability by taking advantage of the sequential nature of agents' decisions. Furthermore—whether in this recursion or in the subsequent counterfactuals—when we solve for one path as a deviation from another, the deviations are always defined between variables of the two paths in the same period. The underlying time-varying fundamentals that are common to both paths, such as migration and trade costs, are therefore 'differenced out.' Our approach thus tractably accommodates arbitrary cross-sectional and temporal variation in these fundamentals, extending the hat algebra method (e.g., Dekle et al., 2007; Caliendo et al., 2019) to dynamic settings with imperfect foresight.

The preceding discussion assumes that beliefs about fundamental productivity are directly observed. Our framework, however, can handle the case where researchers only have data on expectations of some endogenous outcomes (e.g., employment forecasts). The model provides a mapping from any given fundamental belief process to the implied expectations of endogenous outcomes. The researcher's problem is to invert this mapping to identify the unobserved fundamental beliefs (with parameterization) that rationalize the expectation data. Our recursive solution makes this inversion computationally tractable, allowing us to recover both the fundamental beliefs and the full sequence of expected paths simultaneously. We apply this approach in our first quantitative application.

Ex-post counterfactuals. Once the expected paths have been recovered, we can construct counterfactuals by approximating around any of these paths and feeding in appropriate shocks. For instance, the perfect foresight path (the outcome had agents correctly anticipated the productivity shocks from $t = 1$) can be constructed by approximating around the recovered expected path from $t = 1$, $\{\mathbb{E}_1 x_{t'}\}_{t'\geq 1}$, using the belief errors from that period's perspective as the shock: $\{\ln z_{t'} - \mathbb{E}_1 \ln z_{t'}\}_{t'\geq 1}$. This allows us to quantify the impact of imperfect beliefs. In addition to the perfect foresight path, other productivity paths can be used to construct counterfactuals under alternative evolutions of fundamentals or beliefs, following the discussion after Proposition 1.

Why can these recovered expected paths, but not the actual data series, serve as ap-

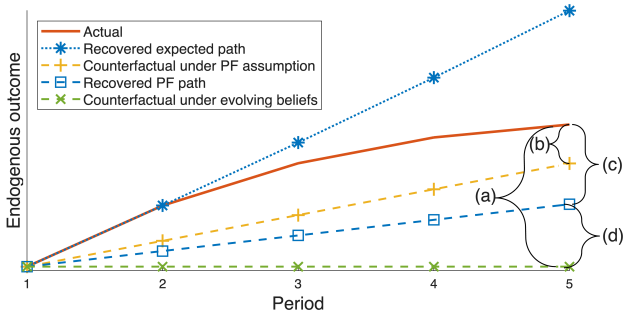

Figure 2: Ex-post Counterfactual Under Alternative Assumptions

Note: Plotted in the figure is an endogenous outcome of interest (e.g. labor share) in the location receiving positive productivity shocks. Red solid line is the actual evolution of the labor share, as in Figure 1. Blue dotted line marked with asterisk is the recovered expected path according to beliefs in the first period, and blue dashed line is recovered perfect foresight path from the first period (PF). Green dashed line illustrates the inferred counterfactual path under evolving beliefs. Yellow dashed line with plus markers is the counterfactual path under the assumption that agents have perfect foresight. Labels (a), (b), (c), and (d) capture the impacts of various mechanisms, see text for explanation.

proximation paths when beliefs are imperfect? As discussed earlier, each expected path constitutes a first-order solution to the dynamic optimization problem agents face in that period. It therefore correctly 'encodes' the influence of the future sequence of migration and trade costs. The actual data series, on the other hand, is a collection of decisions made under different, evolving expectations. It does not satisfy the dynamic optimality conditions of agents, so combining the actual data series with the dynamic optimality conditions of agents leads to internally inconsistent counterfactual outcomes.

Comparison to perfect foresight assumption. We illustrate graphically the importance of accounting for imperfect foresight in ex-post counterfactuals using Figure 2. The red solid line represents the realized evolution, similar to Figure 1. The blue dotted line with circles shows the recovered expected path according to beliefs in the first period, and the blue dashed line with square markers represents the evolution had agents correctly anticipated the productivity path from the first period, constructed as a counterfactual. Suppose the researcher aims to understand the evolution of the economy when productivity remain constant and agents correctly anticipate the productivity. One can solve for this counterfactual by approximating around the expected path, as explained previously.[15]

Let the green dashed line with asterisk markers represent this counterfactual path. The difference between this path and the actual evolution, marked by (a) in the figure, is the causal effect of the productivity shock on the evolution of the economy. Now suppose the researcher assumes instead that the actual evolution of allocations was generated by agents

[15] Alternatively, one can use the recovered perfect foresight path as the approximation path—as long as the deviation in productivity fed into the approximation is defined consistently, different approximation points yield the same first-order result.

with perfect foresight. To obtain the counterfactual path, the researcher takes the observed allocation (red line) as the approximation point and feeds into the system of equations the deviation between the counterfactual (no-growth) productivity path and the actual path. The difference between the resulting counterfactual path and the actual evolution, marked by (b) in the figure, is then taken as the causal effect of the productivity shock.

The two inferred effects, under beliefs (a) and under perfect foresight (b), are different if the beliefs are not perfect foresight. We can decompose this difference into two sources using $(a) - (b) = (c) + (d - b)$. First, component (c) reflects the impact of the 'surprise' component of productivity growth in the observed data. In an environment where there is a discrepancy between actual productivity and agents' belief about productivity, a counterfactual change in productivity necessarily alters this discrepancy, leading to changes in agents' decisions. This component does not arise under the perfect foresight assumption. Second, the difference between (b) and (d) reflects that, by approximating around two different paths while applying the same deviations in productivity, different counterfactual outcomes are obtained. The intuition is that by treating the actual evolution as a perfect foresight path, the researcher implicitly infers incorrect migration costs to rationalize the data.

Even though we motivate the analysis using an example in which agents hold biased beliefs and revise them as new data arrive, the preceding discussion applies whenever the realized fundamentals differ from agents' expected value, as would also occur under rational expectations. We summarize the results in this subsection in the following proposition:

Proposition 2. *i The actual allocations and agents' beliefs about the fundamentals in each period recover the agents' expected paths of outcomes, the perfect foresight paths, and counterfactual paths of outcomes.*

ii Assuming perfect foresight when data are generated by imperfect beliefs yields a different causal effect of shocks to fundamentals.

Proof. See Appendix A.2. □

Throughout this subsection, we have assumed that the researcher observed the data up to T. If the researcher observes data up to period $t' < T$, then before the recursive solution can be applied, the researcher needs to first construct an expected path between t' and T. This can be done analogously to the simulation of ex-ante counterfactuals described in the previous section by taking a stand on the evolution of trade and migration costs.

The general principles of our solution method for both ex-ante and ex-post analyses rely on approximating a stochastic system around a deterministic path, using the deterministic

path and the processes of the fundamentals as inputs. This approach can, therefore, accommodate many extensions, including multiple sectors and other sources of dynamics. We use an extended model with multiple sectors and intermediate goods in the applications.

2.6 Incorporating Uncertainty

We now incorporate the effect of uncertainty into agents' forward-looking decisions. Doing so requires conducting a second-order approximation of the stochastic spatial equilibrium.

We denote by $\bar{x}$ and $\bar{y}$ the outcomes on the deterministic path, $\bar{x}, \bar{y} \in \{\bar{w}_{nt}, \bar{P}_{nt}, \bar{v}_{nt}, \bar{z}_{nt}, ...\}_{n=1}^{N}$; by $\hat{x}$ and $\hat{y}$ the corresponding stochastic deviation terms, $\hat{x}, \hat{y} \in \{\hat{w}_{nt}, \hat{P}_{nt}, \hat{v}_{nt}, \hat{z}_{nt}, ...\}_{n=1}^{N}$. For each outcome after any history $z^{t'}$, we extend the Taylor-expansion of the equilibrium conditions (equations (3) to (8)) around $\bar{x}_{t'}$ from the first to the second order. Taking the expectation from the perspective of period t then delivers a system of equations in both the expected deviations and the covariance of these deviations. In particular, for the labor supply decisions in each $t' = t, t+1, \ldots, T$, this approximation takes the following form:

$$
\begin{aligned}
\mathbb{E}_t \hat{v}_{nt'} &= \frac{\partial U(\bar{c}_{nt'})}{\partial \ln \bar{c}_{nt'}} \mathbb{E}_t \hat{c}_{nt'} + \beta \sum_{i=1}^{N} \bar{\mu}_{nit'} \mathbb{E}_t \hat{v}_{it'+1} && + \frac{1}{2} \sum_{x,y} \frac{\partial^2 F_n}{\partial \bar{x} \partial \bar{y}} \mathbb{E}_t \hat{x}\hat{y} \\
\mathbb{E}_t \hat{\mu}_{nit'} &= \frac{\beta}{\nu} \sum_{m=1}^{N} \left(\mathbb{1}(m = i) - \bar{\mu}_{nmt'} \right) \mathbb{E}_t \hat{v}_{mt'+1} && + \frac{1}{2} \sum_{x,y} \frac{\partial^2 G_{ni}}{\partial \bar{x} \partial \bar{y}} \mathbb{E}_t \hat{x}\hat{y} \qquad (15) \\
\mathbb{E}_t \hat{l}_{nt'+1} &= \underbrace{\sum_{i=1}^{N} \frac{\bar{\mu}_{int'} \bar{l}_{it'}}{\bar{l}_{nt'+1}} \left(\mathbb{E}_t \hat{\mu}_{int'} + \mathbb{E}_t \hat{l}_{it'} \right)}_{\text{first order terms}} && \underbrace{+ \frac{1}{2} \sum_{x,y} \frac{\partial^2 J_n}{\partial \bar{x} \partial \bar{y}} \mathbb{E}_t \hat{x}\hat{y}}_{\text{second order terms}} .
\end{aligned}
$$

This system augments the first-order approximation from Section 2.4 by adding second-order terms that capture the effects of nonlinearity and uncertainty. The functions F_n, G_{ni}, and J_n are the right–hand sides of the equilibrium conditions in Definition 1.

For a concrete example, F_n is defined according to the right-hand side of the value function equation (3): $F_n\left(w_{nt}, P_{nt}, v_{t+1}\right) \equiv U(w_{nt}/P_{nt}) + \nu \ln(\sum_{i=1}^{N} \exp(\beta v_{it+1} - m_{nit})^{1/\nu})$. The first-order terms in the equations above are simply the expected deviations of endogenous variables, weighted by the first derivatives of these functions evaluated at the deterministic path. The second-order terms are the expected cross-products of the stochastic deviations, $\mathbb{E}_t[\hat{x}\hat{y}]$, weighted by their corresponding second-derivative coefficients $\frac{\partial^2 F_n}{\partial \bar{x} \partial \bar{y}}$. In Appendix A.3, we show that these cross-derivative coefficients can be characterized as closed-form functions of outcomes on the deterministic path (e.g., trade and migration shares). We provide similar expressions for trade and labor demand equilibrium conditions, thereby

establishing the following proposition:

Proposition 3. *Given the labor allocations in period t, agents' expectation in t about the deviations in exogenous fundamentals and endogenous variables in $t' \geq t$ from a deterministic path* $\left\{\mathbb{E}_t\hat{z}_{it'}, \mathbb{E}_t\hat{v}_{it'}, \mathbb{E}_t\hat{\mu}_{nit'}, \mathbb{E}_t\hat{l}_{it'}, \mathbb{E}_t\hat{w}_{it'}, \mathbb{E}_t\hat{\lambda}_{nit'}, \mathbb{E}_t\hat{P}_{it'}\right\}_{i=1,n=1,t'=t}^{N,N,T}$ *and the covariance of these deviations* $\left\{e.g., \mathbb{E}_t\hat{z}_{it'}\hat{z}_{nt'}, \mathbb{E}_t\hat{v}_{it'}\hat{v}_{nt'}, \mathbb{E}_t\hat{w}_{it'}\hat{w}_{nt'}, \mathbb{E}_t\hat{w}_{it'}\hat{P}_{nt'}\right\}$ *solve a system of linear equations (in the order of $N^2 \times (T-t)$), with their coefficients being closed-form expressions of outcomes on the deterministic path.*

Proof. See Appendix A.3. □

This proposition shows that the system of equations in Proposition 1, when appended with second-order terms, characterizes the solution with second-order accuracy. The inputs are now the equilibrium outcome in the deterministic transitional path, and both the first and second moments of the deviations in stochastic productivities, i.e. $\mathbb{E}_t\hat{z}_{nt'}$ and $\mathbb{E}_t\hat{z}_{it'}\hat{z}_{nt'}$. The unknowns now include not only the expected deviations but also the covariance terms involving endogenous variables (e.g., $\mathbb{E}_t[\hat{w}_{it'}\,\hat{P}_{nt'}]$). To see how these second-order terms matter, consider a rational expectations equilibrium with $\mathbb{E}_t\hat{z}_{it'} = 0$. In first-order approximations, the expected deviations of all endogenous variables (e.g. $\mathbb{E}_t\hat{v}_{nt'}$) are zero. In second-order approximations, the presence of the variance-covariance of fundamental productivities ($\mathbb{E}_t\hat{z}_{it'}\hat{z}_{nt'}$) means $\mathbb{E}_t\hat{v}_{nt'} = 0$ is not a solution, reflecting the impacts of uncertainty.

Recall also that in first-order approximations described in Proposition 1, we can simply invert the system of equations to obtain a solution. Here, the presence of the (unknown) covariance terms between endogenous outcomes makes the system under-ranked. To solve the system with second-order accuracy, we devise a two-step procedure.

Suppose we are solving agents' decision in period t. First, we compute the endogenous covariance terms, $\{\mathbb{E}_t\hat{x}_{t'}\hat{y}_{t'}\}_{t'\geq t}$, using simulations from the *first-order* solution. Specifically, we simulate S paths of future productivity according to beliefs at t, and for each simulated path $s = 1, 2, ..., S$ we apply Proposition 1 recursively to obtain the outcomes of the economy, denoted by $\{\hat{x}_{t'}(s)\}_{t'\geq t}$ or $\{\hat{y}_{t'}(s)\}_{t'\geq t}$. We then approximate $\mathbb{E}_t\hat{x}_{t'}\hat{y}_{t'} \approx \frac{1}{S}\sum_{s=1}^{S}\hat{x}_{t'}(s)\hat{y}_{t'}(s)$. Second, we plug these simulated covariance terms back into the second-order system from Proposition 3 and solve the now-determined linear system for the remaining unknowns—the expected deviations, $\{\mathbb{E}_t\,\hat{x}_{t'}\}_{t'\geq t}$.

The resulting solution is accurate up to the second order. The intuition is as follows:

$$\begin{aligned}\mathbb{E}_t\hat{x}_{t'}\hat{y}_{t'} &= \frac{1}{S}\sum_{s=1}^{S}\left[\hat{x}_{t'}(s)+o\left(\hat{x}_{t'}(s)\right)\right]\cdot\left[\hat{y}_{t'}(s)+o\left(\hat{y}_{t'}(s)\right)\right]+O(1/\sqrt{S})\\ &= \frac{1}{S}\sum_{s=1}^{S}\Big[\hat{x}_{t'}(s)\hat{y}_{t'}(s)+\underbrace{o\left(\hat{x}_{t'}(s)\right)o\left(\hat{y}_{t'}(s)\right)+\hat{y}_{t'}(s)o\left(\hat{x}_{t'}(s)\right)+\hat{x}_{t'}(s)o\left(\hat{y}_{t'}(s)\right)}_{\text{third-order terms}}\Big]+O(1/\sqrt{S})\end{aligned}$$

where $o\left(\hat{x}_{t'}(s)\right)$ and $o\left(\hat{y}_{t'}(s)\right)$ represent the second-order errors in solving for $\hat{x}_{t'}(s)$ and $\hat{y}_{t'}(s)$ using Proposition 1, and $O(1/\sqrt{S})$ denotes the simulation error for the covariance. For a large enough S, the simulation error becomes arbitrarily small.[16] The remaining approximation error arises from terms that are products of a second-order term with a first or higher-order term (e.g. $\hat{x}_{t'}(s)o\left(\hat{y}_{t'}(s)\right)$). These terms are therefore of the third order. Hence, by substituting the simulated covariances into the second-order system and solving for the expected deviations, we arrive at a solution with second-order accuracy.[17, 18]

We can use Proposition 3 recursively to obtain the evolution of the economy under any given path of fundamentals. It simulates the transitional dynamics of the economy in ex-ante analyses; it can also be applied to each of the steps in the recursive solution described in Section 2.5 to obtain second-order solution in ex-post counterfactuals when the data are generated by agents holding uncertain beliefs.

Comparison to state-space methods. Our approach relates to local perturbation methods, which typically operate in state space and approximate decision rules around a steady state using numerical Jacobians and Hessians (see e.g., Kim et al., 2008; Schmitt-Grohé and Uribe, 2004).[19] This method has been extended to incorporate heterogeneous agents or

[16] A concern is that the number of simulations necessary for approximation $\mathbb{E}_t\hat{x}_{t'}\hat{y}_{t'}$ increases with the number of regions. Note, however, that by using Monte Carlo integration, our algorithm has a convergence rate of $O(1/\sqrt{S})$, independent of the dimension of the state space.

[17] Appendix A.5 verifies in a deterministic setting the accuracy of the second-order solution for large shocks. Note also that when comparing second- and first-order solutions, the difference could capture both the effects of uncertainty and deterministic nonlinearity. In the application, we design the baseline to isolate the effect of uncertainty.

[18] An alternative to our simulation-based approach is to express the covariance terms analytically as functions of the second-moments of the fundamental shocks—see Fan and Luo (2025) and Kleinman et al. (2025) for development of this approach in static trade models. However, in our dynamic multi-sector trade model with intermediate goods and history dependence, any shock can affect the covariance between any two future endogenous outcomes. Deriving these relationships requires computing the Hessian of every outcome with respect to every past shock in all locations, which is analytically and computationally intractable.

[19] In standard perturbation methods, second-order terms can accumulate and lead to explosive paths. To address this, Kim et al. (2008) propose 'pruning' higher-order terms. Our approach avoids such accumulation by computing the relevant covariance terms using first-order simulations.

multiple regions (e.g., Bhandari et al., 2023; Bilal, 2023). There are two main differences between our approach and standard perturbation. First, we operate in sequence rather than state space. Second, we approximate around a given path instead of the steady state.

Our approach is well suited to studying the transition dynamics of an economy under a pre-specified productivity path and belief structure—many of the applications referenced in the introduction fall into this category—offering two key advantages for such questions. First, focusing on sequences rather than states accommodates time-varying fundamentals, flexible productivity processes, and history-dependent beliefs without increasing computational complexity. In contrast, the dimensionality of state-space methods grows rapidly as these features become more flexible.[20] Second, we leverage the fact that, for a broad class of dynamic spatial models, the derivatives (first and higher) of outcomes with respect to prices and other endogenous outcomes are closed-form expressions of outcomes on the approximation path, circumventing the need for numerical derivatives in state-space methods.

By contrast, a state-space representation of the model is better suited for characterizing the stochastic stationary equilibrium. In some applications, researchers are interested in the model's second moments in the stochastic stationary equilibrium—for instance, the business cycle properties of the spatial economy or the volatility and co-movement of regional outcomes. For this type of exercise, solving a time-invariant policy function via a state-space method can be computationally efficient.

Finally, we note that in environments with uncertainty, agents may also have the option to self-insure through savings. For shocks that are persistent, slowly unfolding, and aggregate in nature, such as the climate change shock examined in our quantification, this form of insurance may provide limited benefit. Exploring the feasibility of incorporating savings into our simulation-based approach in other contexts—for example, by extending the cross-derivatives to saving decisions—represents an interesting direction for future research.

[20] For illustration, if the productivity process is AR (1) and common across locations, the state space of the model in t is (l_t, z_t), with a dimension of $N+1$; if productivity processes are location-specific, the dimension is $2 \times N$; if either the productivity or agents' beliefs depend on x period lagged outcomes, the dimension increases to $2 \times N \times x$. The dimension expands further with additional features. For example, if migration and trade costs vary flexibly over time—an important feature in ex-post applications of spatial models—then the state-space representation of the model differs across periods, requiring separate approximations for each period. In the extension of the model with heterogeneous beliefs, higher-order beliefs become a part of the state. Using a state-space representation, most existing studies of dynamic spatial models with uncertainty adopt relatively low dimension productivity processes and do not incorporate these additional features.

2.7 Welfare Measurement without Perfect Foresight

A key goal in dynamic spatial models is to evaluate how a shock affects not only allocations but also agents' welfare, which can be measured using the average (across idiosyncratic tastes) location value, v_{nt}. Under perfect foresight, this measurement is straightforward in both ex-ante and ex-post analyses. In ex-ante analysis, a researcher assumes a path for all fundamentals, and v_{nt} can be solved using the deterministic version of equation (3). In ex-post analysis, while the fundamentals are unknown, their influence on welfare is summarized by observed real wages and migration rates. Intuitively, the sequence of real wages reflects the value of living in a location if agents are forced to stay, while migration choices reveal the option value of moving. Welfare can therefore be measured using these observed outcomes without knowledge of the underlying fundamentals (see, e.g., Caliendo et al., 2019).[21]

In environments without perfect foresight, we must distinguish between three notions of welfare. The first is **ex-ante welfare**: the expected value at each period according to agents' beliefs, which could be evolving and uncertain. The second is **perfect foresight welfare**: the value agents would have experienced had they correctly anticipated the realized path of fundamentals. The third is **realized welfare**: the value that corresponds to the actual path of allocations, where agents make decisions based on beliefs that they subsequently revise as new information arrives.

In the previous subsections, we described how to construct the expected and perfect foresight paths of the economy with first- and second-order accuracy in both ex-ante and ex-post analyses. Leveraging these paths, we can compute the first two notions of welfare. Specifically, from equations (3) and (4), we obtain:

$$\mathbb{E}_t v_{nt} = U(c_{nt}) - \nu \ln \mu_{nnt} + \beta \mathbb{E}_t v_{nt+1}. \tag{16}$$

By substituting forward this equation using the sequence of equilibrium outcomes on the expected path from t, we can recover ex-ante welfare from the perspective of period t: $\mathbb{E}_t v_{nt}$, $\mathbb{E}_t v_{nt+1}$, $\mathbb{E}_t v_{nt+2}$, $\mathbb{E}_t v_{nt+3}$, ... Perfect foresight welfare can be constructed analogously using equilibrium outcomes along the counterfactual perfect foresight path.

Measuring the third notion, realized welfare—arguably the most relevant in many applications of dynamic spatial models—is more challenging. Forward substitution fails because agents' decisions on the actual path are shaped by beliefs that are subsequently revised; the

[21] Formally, the value of a location under perfect foresight is $\bar{v}_{nt} = U(\bar{c}_{nt}) - \nu \ln \bar{\mu}_{nnt} + \beta \bar{v}_{nt+1}$. Substituting this equation forward recursively expresses $\bar{v}_{nt}$ as a discounted sum of $U(\bar{c}_{nt}) - \nu \ln \bar{\mu}_{nnt}$.

realized values at t and $t+1$ are therefore not linked by a simple forward-looking Bellman equation. We show, however, that realized welfare can be expressed as the sum of ex-ante welfare and a correction term reflecting the 'surprise' in fundamentals. Both components can be recovered from allocations without knowledge of the underlying fundamentals.

Formally, let $\tilde{v}_{nt}$ be the realized welfare in location n at time t and let $\mathbb{E}_t v_{nt+1}$ be the expected value at t about the value at $t+1$. These two values are related as follows:

$$\begin{aligned}
\tilde{v}_{nt} &= U(\tilde{c}_{nt}) + \mathbb{E}_\epsilon \sum_{i=1}^{N} \mathbb{1}_{\text{i is chosen}}[\beta \tilde{v}_{it+1} - m_{nit} + \epsilon_{it}] \\
&= U(\tilde{c}_{nt}) + \mathbb{E}_\epsilon \sum_{i=1}^{N} \mathbb{1}_{\text{i is chosen}}[\beta \mathbb{E}_t v_{it+1} - m_{nit} + \epsilon_{it}] + \mathbb{E}_\epsilon \sum_{i=1}^{N} \mathbb{1}_{\text{i is chosen}}[\beta \tilde{v}_{it+1} - \beta \mathbb{E}_t v_{it+1}] \\
&= U(\tilde{c}_{nt}) - \nu \ln(\tilde{\mu}_{nnt}) + \beta \mathbb{E}_t v_{nt+1} + \sum_{i=1}^{N} \tilde{\mu}_{nit}[\beta \tilde{v}_{it+1} - \beta \mathbb{E}_t v_{it+1}]. \qquad (17)
\end{aligned}$$

The first equality writes the actual location value as the sum of the flow utility, which depends on the realized consumption $\tilde{c}_{nt}$, and the continuation value. The latter depends on agents' *actual choice* given their beliefs at t (a function of ϵ_{it} summarized by the indicator function), where the expectation $\mathbb{E}_\epsilon$ denotes integration over ϵ_{it}. The second line adds and subtracts $\beta \mathbb{E}_t v_{it+1}$ to the continuation value. The third equality uses the property that the actual migration decision $\tilde{\mu}_{nit}$ is made according to the *expected* values, which implies $\mathbb{E}_\epsilon \sum_{i=1}^{N} \mathbb{1}_{\text{i is chosen}}[\beta \mathbb{E}_t v_{it+1} - m_{nit} + \epsilon_{it}] = [\beta \mathbb{E}_t v_{nt+1} - \nu \ln(\tilde{\mu}_{nnt})]$.

Equation (17) thus expresses realized welfare in period t as a function of the realized allocation $(\tilde{c}_{nt}, \tilde{\mu}_{nnt})$, the expected location value in $t+1$ based on beliefs in t $(\mathbb{E}_t v_{n,t+1})$, and the realized value in $t+1$ $(\tilde{v}_{i,t+1})$. Because we can calculate $\mathbb{E}_t v_{nt+1}$ as described above using the expected path from t, this equation forms a backward recursion. Starting from the last period of the model, or from any period in which agents gain perfect foresight in infinite horizon models, we can iterate on this equation backward to obtain the realized welfare in every location and period—all without needing to back out fundamentals such as trade and migration costs.

Accuracy. Because equations (16) and (17) are exact, their accuracy depends solely on the accuracy of the paths used as inputs. Since our methods can construct these paths with second-order accuracy, the resulting welfare measures are also second-order accurate.

In setting up the problem, we have not distinguished whether the researcher is simulating the model forward taking an assumed path of fundamentals as given or conducting ex-post analyses taking the observed data as given. In both types of analyses, as long as the expected,

perfect foresight, and the realized paths are constructed, all three welfare measures can be evaluated. We summarize the results in this section in the following proposition:

Proposition 4. *In both ex-ante and ex-post analyses, the ex-ante, perfect foresight, and realized welfare of locations can be recovered from the corresponding allocation paths without knowledge of the underlying fundamentals, up to the second order.*

2.8 Summary and Extension to Heterogeneous Beliefs

In this section, we have developed a generic dynamic stochastic spatial model with flexible beliefs. We derived an approximate solution for this model around a deterministic path. We then developed a methodology to study the effects of evolving beliefs on the evolution of the economy for both ex-ante (taking model primitives as given) and ex-post (taking the observed data evolution as given) analyses. For both types of analyses, the impacts of uncertainty can be incorporated through second-order approximations, and relevant welfare measures can be constructed.

We have so far focused on the case in which all agents share a common belief f. In environments with stochastic fundamentals, agents might disagree about the path of future fundamentals. In such cases, agents must take into account not only their own beliefs but also the beliefs of other agents, as their decisions, too, affect equilibrium outcomes. This leads to a well-known curse-of-dimensionality due to higher-order beliefs (Townsend, 1983). We relegate to Appendix A.4 an extension of the model to incorporate heterogeneous beliefs among groups of agents. We retain tractability by proposing three alternative assumptions on the structure of higher-order beliefs, under which we characterize the equilibrium conditions and develop solution algorithms.[22] We describe how the solution for heterogeneous beliefs can be extended to accommodate uncertainty in beliefs as described in the previous section. We further demonstrate how heterogeneous beliefs can be incorporated into the analyses of ex-post counterfactuals—i.e., when researchers observe the actual data governed by groups of agents holding different beliefs.

We now turn to quantitative analyses, applying our methodologies to two settings: China's productivity growth and climate change. Through these two applications, we explore the roles of evolving, uncertain, and heterogeneous beliefs in dynamic spatial models.

[22] Huo and Takayama (2024) shows that in rational-expectation models with dispersed information, when signals follow finite ARMA processes, the equilibrium outcomes admit a finite-state representation. Our setting differs from that in Huo and Takayama (2024) as we allow for time-varying fundamentals and flexible beliefs.

3 Ex-post Application: The China Shock

The first application is an ex-post analysis of the spatial effects of the "China shock" in the presence of evolving beliefs. We study the impact of rapid productivity growth in China on manufacturing employment and welfare in the U.S. in a setting without perfect foresight.

Figure 3 illustrates the relevance of departing from perfect foresight in the context of the China shock. The figure shows that the U.S. manufacturing employment share experienced a steep and persistent decline after 1990, before flattening out around 2008. Crucially, the speed and persistence of this decline were not fully anticipated. The dotted lines plot 10-year employment projections from the Bureau of Labor Statistics (BLS), which were initially overly optimistic but were revised downward over time, eventually converging to the actual path. Coinciding with this manufacturing employment decline was the rapid productivity growth in China. To the extent that agents in the economy also failed to foresee the severity of this decline—whether because they did not anticipate China's productivity growth or other economic changes—taking their evolving beliefs into account is crucial for studying the impact of the China shock.

We model China's productivity dynamics as a catch-up process with that of the U.S., denoted by $g(z_{t+1}|z^t)$. We allow agents to form beliefs $f(z_{t+1}|z^t)$ about $g(z_{t+1}|z^t)$. We then use our recursive solution to recover this unobserved belief process by requiring that the model, when fed these beliefs, generates the manufacturing employment projections made by the BLS at each point in time.

3.1 Setup and Data

We take the model to a geographic setting that includes the 50 U.S. states, 27 other countries/regions, and a synthetic rest of the world. We also extend the baseline model to incorporate multiple sectors and intermediate goods (see Appendix B.1 for details). Since imports from China consist primarily of manufacturing goods, we include two sectors—manufacturing and non-manufacturing. We allow for migration across sectors and states within the United States, but not internationally. In addition, we incorporate trade across U.S. locations as well as bilateral international trade between individual U.S. states and other countries.

Our main data span the period 2000–2016. We construct yearly gross mobility flows across U.S. states and sectors using job-to-job flow data from the U.S. Census Bureau's

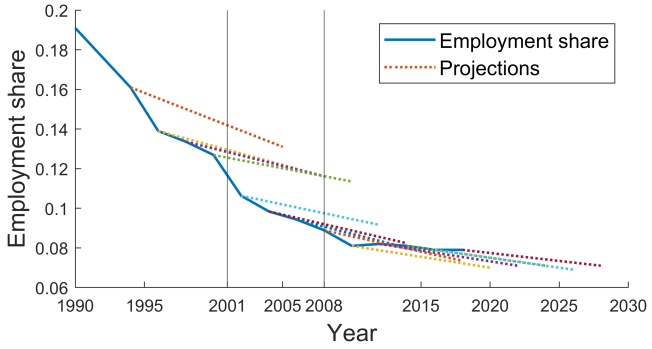

Figure 3: U.S. Manufacturing Employment Share

Note: This figure presents the manufacturing employment share in the United States during 1990-2018 (solid line) and 10-year projections by the Bureau of Labor Statistics (dotted lines).

LEHD program.[23] We obtain bilateral manufacturing trade flows across U.S. states from the Commodity Flow Survey (CFS), reported every five years over the period 2002–2017. We interpolate between reporting years using data from the nearest available year. We approximate trade in agriculture and services—together they form the non-manufacturing sector in the model—using an estimated gravity equation. International trade data are obtained from the World Input-Output Database (WIOD). We allocate bilateral international trade to individual U.S. states based on each state's gross output share in a given industry. We also obtain input shares from WIOD, and data on gross output and value added from the Bureau of Economic Analysis (BEA) for the U.S. states and from WIOD for other countries. We set the migration elasticity to $1/\nu = 0.187$, following Caliendo et al. (2019), the trade elasticity to $\theta = 4.55$, based on Caliendo and Parro (2015), and a discount factor $\beta = 0.96$. Appendix B.2 provides further details on data sources and construction.

Our key modeling assumption is that agents have perfect foresight about all fundamentals except for China's manufacturing productivity, with beliefs disciplined by the BLS employment projections shown in Figure 3. Two aspects of this choice warrant discussion.

First, we focus on beliefs about China's productivity as the source of the forecast errors. The very term 'China shock,' common in the literature, suggests the surprise was external to the U.S. economy. This view is supported by the extensive research linking U.S. manufacturing's decline to trade with China (Autor et al., 2013; Pierce and Schott, 2016; Handley and Limão, 2017) and is consistent with observed productivity dynamics, as we show below: the U.S.-China productivity gap closed primarily due to rapid growth in China, not

[23] This dataset is not suitable for measuring stayers in a location because it omits individuals who stay in the same job. We calculate stayers in each state-sector cell as the difference between total employment and outflows.

an unexpected slowdown or other structural breakdown in the U.S. While other sources of misperception are possible, we view this as the most plausible channel.[24]

Second, we model the forecast errors as arising from biased beliefs—where the perceived mean of the process is incorrect—rather than from pure uncertainty, where the mean is correct but outcomes are noisy. The persistent, one-sided nature of the BLS forecast errors in Figure 3 is more consistent with a systematically biased first moment than with random shocks from a correctly centered distribution.

3.2 Productivity Catch-up and Evolving Beliefs

To guide subsequent analysis, we first estimate the data generating process for China's productivity catch-up. Using a model inversion based on domestic trade shares, wages, and prices, we construct a time series of manufacturing productivity, z_t^{nj}, for both China and the United States over 1990-2016 (See Appendix B.2.2 for details).

Figure 4a plots the resulting log difference in manufacturing productivity between China and the U.S. As the figure shows, after a decade of modest progress, China's productivity catch-up accelerates dramatically around 2000, following an S-shaped trajectory. Motivated by the S-shaped pattern in relative productivity after 2000, we model the catch-up process as a logistic function:

$$\ln(z_t^{CN,M}) - \ln(z_t^{US,M}) = \alpha_3 \cdot \frac{\exp\left(\alpha_1(t-1999)\right)}{\exp\left(\alpha_1(t-1999)\right) + \alpha_2}, \tag{18}$$

where the coefficients α_1, α_2, and α_3 govern the catch-up speed, initial gap, and average productivity of China relative to the U.S., respectively. Estimating this specification on post-2000 data yields parameter values of $\alpha_1 = -0.22$, $\alpha_2 = 0.11$, and $\alpha_3 = -2.4$. The solid line in Figure 4a shows that the estimated function fits well the observed catch-up process, which further justifies modeling beliefs as learning about a deterministic trend, rather than as a response to a series of negative U.S. productivity shocks.

We treat the estimated catch-up model as the true data generating process, $g(z_{t+1}|z^t)$, for China's relative productivity to the U.S. We use it to project the path of China's absolute productivity level beyond 2016, assuming U.S. manufacturing productivity remains constant after 2016. The resulting estimated path, shown by the solid line in Figure 4b, has China's

[24] We also note that in a two-country model, it is the relative manufacturing productivity that matters for labor allocation. In such an environment, whether the misperception is about China's productivity, U.S. productivity, or other manufacturing-specific factors (e.g., amenities) has little impact on the counterfactuals.

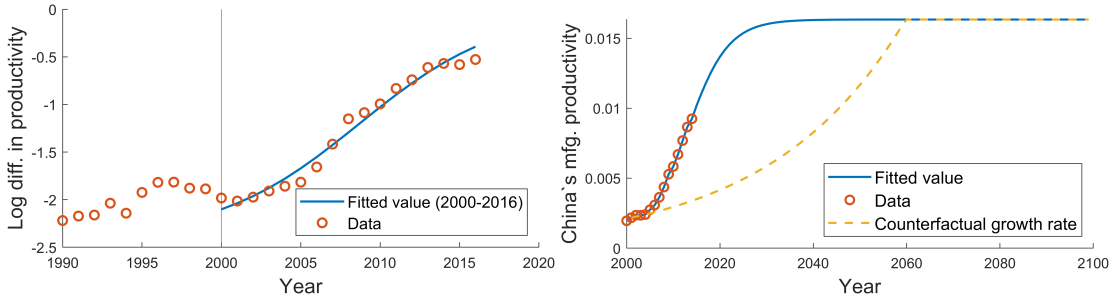

(a) CHN–U.S. Log Productivity Difference

(b) Fitted Productivity for China

Figure 4: Measured Manufacturing Productivity in the United States and China

Note: Panel (a) plots the log difference in manufacturing productivity between China and the U.S. from 1990-2016. The solid line shows the fit of a logistic function estimated on post-2000 data. Panel (b) shows the level of China's manufacturing productivity implied by this fitted model (solid line) and the counterfactual 'no-China-shock' path assuming a slower growth rate (dashed line; see Section 3.3 for details).

productivity continuing to grow until it converges with the U.S. level around 2040.

Next, we recover agents' evolving beliefs about this catch-up process. We assume agents form beliefs about the catch-up speed, captured by α_1 (while knowing α_2 and α_3), and update these beliefs as new data arrive. Our goal is to recover agents' perceived value of α_1—and by implication, their belief about the fundamental productivity path, $f(z_{t+1} \mid z_t)$—for each year from 2001 to 2006.

The method developed in Section 2.5 provides a theory-consistent mapping from any productivity belief to an expected employment path. We search for the belief parameter, α_1, that makes the model's implied employment forecast match the corresponding ten-year BLS projection from Figure 3. As we show in Appendix B.2.3, this reveals that agents were initially pessimistic about China's catch-up but gradually revised their beliefs, nearly fully aligning them with the true process by 2006.

We recover these beliefs using the recursive solution, which decomposes the joint identification of beliefs across all periods into a sequence of separate identification problems, one for each period. In each step of the backward recursion, we identify one unknown parameter (α_1 for that year) to match one data moment (the BLS forecast from that year). Because the model-implied future manufacturing employment share is monotonic in the perceived productivity path, this belief parameter is identified. This, in turn, allows for the unique recovery of the full expected path for that year. Since this holds for every step, this process identifies the entire sequence of beliefs and expected paths.

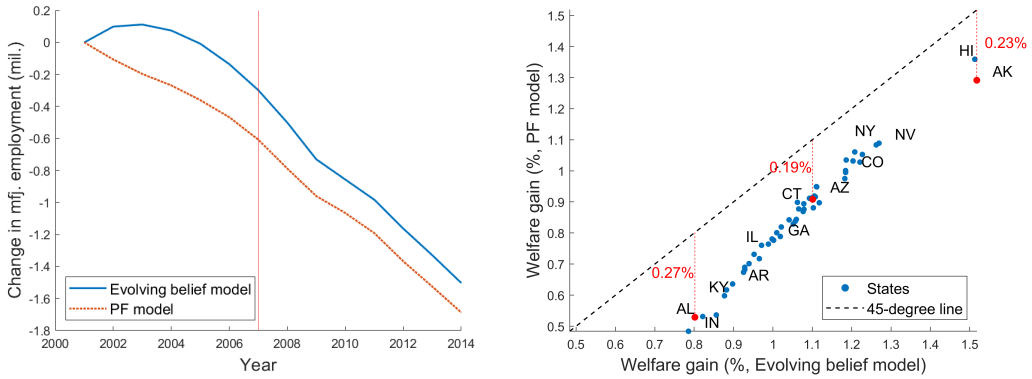

(a) Employment Effect (b) Welfare Effect

Figure 5: Employment and Welfare Effect of the China Shock

Note: The figure compares the inferred effects of the China shock under evolving beliefs and perfect foresight (PF). Panel (a) shows the change in U.S. manufacturing employment relative to the no-shock counterfactual. Panel (b) plots the welfare gains for each state in the model with evolving beliefs (x-axis) against the gains under the PF model (y-axis). States below the 45-degree line experience larger welfare gains under evolving beliefs. The annotated numbers on Panel (b) are the percentage point differences between inferred gains by the model with evolving beliefs and the one with perfect foresight.

3.3 The Impact of the China Shock under Evolving Beliefs

With the true productivity process and the recovered belief process in hand, we apply our methodology to study the impact of the China shock. We interpret the 'shock' as the rapid, and partially unanticipated, productivity catch-up of China. Accordingly, we define the no-shock counterfactual as a scenario where China's manufacturing productivity grows at a slower, perfectly anticipated rate. Specifically, we assume a 3.5% annual growth rate, which implies China's productivity converges to the U.S. level in 2060, roughly 20 years later than in the data (see Figure 4b).[25]

We address two questions under this setup. First, what was the total effect of the China shock on the U.S. economy, incorporating both the change in productivity and agents' evolving beliefs about it? We measure this as the difference between the realized path and the no-shock counterfactual. Second, how do these inferred effects differ when we account for evolving beliefs versus when perfect foresight is assumed?

Figure 5a plots the inferred impact on U.S. manufacturing employment over 2000-2014. The blue solid line shows the employment decline in the model with evolving beliefs, relative to the no-shock counterfactual. The red dotted line shows the same comparison but under the perfect foresight (PF) assumption. If agents had perfect foresight, they immediately started reallocating out of manufacturing. In contrast, if agents held evolving beliefs, the reallocation was slower. Being initially pessimistic about China's catch-up, agents only

[25] Alternative counterfactual productivity paths are presented in the Appendix B.2.3 for comparison.

began to reallocate as they updated their beliefs over time. This initial divergence has persistent effects; the cumulative job loss inferred by the model with evolving beliefs is more than 10% smaller than that inferred under perfect foresight by 2014, eight years after agents' beliefs converged to the true process.[26]

Figure 5b reports the welfare effect of the China shock on U.S. states under evolving beliefs and perfect foresight. Welfare is defined based on the average location values in 2000. For the case with evolving beliefs, we construct the realized welfare as in Section 2.7. We find that the model with evolving beliefs infers larger gains than the model with perfect foresight. This is visible in the plot, where all states lie below the 45-degree line. The difference is sizable; for example, Alabama experiences approximately 0.8% welfare gains in the model with evolving beliefs, 0.27 percentage points higher than the 0.53% gain under perfect foresight. In the aggregate, the gains under evolving beliefs (1.11%) are about one-third larger than under perfect foresight (0.81%).

This finding may seem counter-intuitive, as one might expect perfect foresight to improve welfare. Note, however, that this exercise is not about the value of information for agents, but rather about how a researcher's assumptions shape the inferred welfare effects. The core mechanism operates through the inferred migration costs. In the model with evolving beliefs, the observed slow reallocation out of manufacturing is rationalized by agents' initial pessimism about China's productivity. A model that assumes perfect foresight must instead rationalize this same slow reallocation with higher migration costs. Indeed, we show in Appendix B.2.5 that under perfect foresight, to fit the manufacturing employment, the implicit cost of moving out of manufacturing would need to be about 50% higher.[27]

This difference impacts the inferred welfare effects. Because our model infers lower migration costs, lower adjustment costs result in higher net welfare gains from the China shock. An additional channel related to the terms of trade reinforces this result. The slower transition out of manufacturing implied by evolving beliefs leads to an improvement in the U.S. terms of trade, further increasing the calculated welfare gains.

[26] Appendix Figure B.2a presents the long-run dynamics of manufacturing employment over 2000-2100. In both cases, the manufacturing decline due to the shock continues to intensify after 2014 as the gap between the actual and counterfactual productivity paths continues to widen. This decline eventually reverses since both the actual and counterfactual productivity paths converge to the same level (see Figure 4b).

[27] Specifically, in Appendix B.2.5, we multiply the manufacturing-to-non-manufacturing migration costs by $(1 + x)$ over 2001-2007 from the perfect foresight path recovered by the model with evolving beliefs. We choose x so that the new path has the same manufacturing employment in 2007 as in the data. We find $x \approx 0.5$.

4 Ex-ante Application: Climate Change

Our second application studies the effects of climate change on welfare and the spatial distribution of economic activity in the U.S. Since the full impact of climate change is yet to unfold, we use this ex-ante setting to illustrate how our framework can handle two departures from perfect foresight that are especially relevant in this context: uncertainty about the extent of future climate change and heterogeneous beliefs about whether it will occur at all.

4.1 Setup and Data

As in the previous application, our geography includes the 50 U.S. states, 27 other countries/regions, and a rest-of-the-world aggregate. We incorporate three sectors—agriculture, manufacturing, and non-tradable services—and intermediate goods. We study the impact of rising temperatures over the period 2014–2100 under different belief assumptions. Our strategy is to use a path of fundamentals that reflects the predicted effect of climate change on productivity and compute the corresponding deterministic path. We then apply Proposition 3 to study the effect of uncertainty about the path of climate change and the result summarized in Section 2.8 to quantify the role of heterogeneous beliefs about climate change. Since one of the key focuses of our quantitative analysis is uncertainty, agents' risk attitude plays an important role. We specify a constant risk aversion (CRRA) flow utility from consumption with a risk-aversion parameter $\sigma = 3$. We treat each period as two years and choose β so the annual discount rate is 4%.

This exercise requires data on gross labor mobility flows across states and sectors; bilateral trade between U.S. states and foreign countries, input-output tables of each country, and U.S. state level sectoral compositions. Since most data sources were described in the previous section, we defer the details on these variables to Appendix B. Here, we outline the key new inputs for this application: the temperature projections and damage functions that map temperature changes into productivity shocks.

Future temperature rise and productivity damage. To construct a path of fundamentals that reflects the impact of climate change, we combine temperature projection data with existing estimates of temperature's effect on productivity.

Our temperature projection data are from the Climate Impact Lab, which provides forecasts under different carbon emission scenarios (RCPs). Figure 6a depicts these global

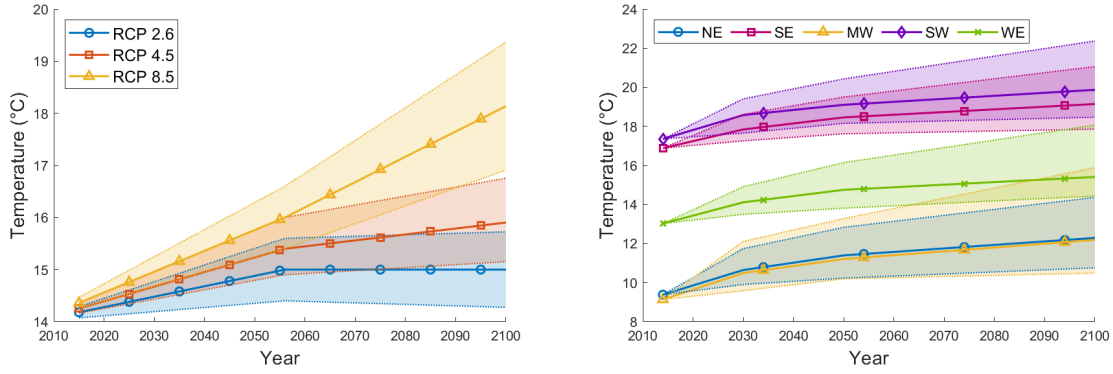

(a) Global Temperature Projection

(b) Local Temperature Projection

Figure 6: Temperature Projection

Note: Panel (a) presents the global average surface temperature change relative to 1986–2005. The graph displays projected time series along with uncertainty estimates (5–95% confidence intervals), depicted by solid lines and shaded areas, for three scenarios: RCP 2.6 (blue – low-emission, strong mitigation), RCP 4.5 (red – medium-emission, moderate mitigation), and RCP 8.5 (yellow – high-emission, business-as-usual). Source: AR5 Synthesis Report IPCC 2014. Panel (b) shows projected temperature changes under RCP 4.5 across broad U.S. regions: Northeast (NE), Southwest (SW), West (WE), Southeast (SE), and Midwest (MW). Solid lines indicate the median forecasts, while shaded areas represent the 5th and 95th percentiles.

projections. The figure highlights two sources of uncertainty: first, uncertainty about which emissions path the world will follow (e.g., RCP 2.6 vs. 8.5), and second, for any given path, considerable uncertainty among the scientific community about the resulting temperature change (the shaded confidence bands), which expands over longer horizons.

We focus on the uncertainty within a single, intermediate scenario: RCP 4.5. Figure 6b shows the 5th, 50th (median), and 95th percentile temperature paths for five broad U.S. regions under this scenario. Two features are notable. First, the wide confidence bands reflect the significant uncertainty about the future temperature path. Second, the southern regions (SW and SE) are consistently hotter than the rest of the country. Given the non-linear relationship between temperature and productivity, this implies that future warming will have a disproportionately negative impact on the South.

We specify the following damage function to link temperature changes to productivity:

$$\begin{aligned} \ln(z_t^{nj}) &= \delta^{nj} + \alpha_1^j \bar{C}_{nt} + \alpha_2^j (\bar{C}_{nt})^2 \\ \implies \Delta \ln(z_t^{nj}) &= \alpha_1^j \Delta \bar{C}_{nt} + \alpha_2^j \Delta (\bar{C}_{nt})^2 \end{aligned} \tag{19}$$

where $\Delta \ln(z_t^{nj})$ is the change in log productivity in location n and industry j, and $\bar{C}_{nt}$ is the average annual temperature. The coefficients α_1^j and α_2^j capture the impacts of temperature. Given a temperature path, this equation yields the corresponding productivity path.

We calibrate coefficients α_1^j and α_2^j using the estimates from Cruz (2023), who estimate

Table 1: Quadratic Impact of Temperature on Productivity

	Agriculture	Manufacturing	Service
α_1^j	2.998	1.339	1.292
	(1.268)	(1.344)	(1.051)
α_2^j	-0.143	-0.0615	-0.0485
	(0.0368)	(0.0472)	(0.0405)

Note: Estimates and standard errors (in parentheses) based on Cruz (2023).

the long-run effect of climate change on sectoral productivity using global data.[28]

Table 1 reports the estimates. Although the coefficients differ across sectors, all sectors show an inverted U relationship between temperature and productivity, which means that the marginal damage from an additional degree of warming is greater in already hot locations. Indeed, in our simulation, the U.S. South will experience larger productivity decline than other U.S. regions (see Appendix Figure B.3). The standard errors for the coefficient estimates are large, reflecting the inherent difficulty of projecting the economic toll of climate change. We will incorporate the uncertainty surrounding both the temperature path and these damage function coefficients into agents' beliefs.[29]

Constructing the baseline deterministic path. With the initial allocation data and the damage function defined, we can now construct the baseline deterministic path that will serve as the foundation for our counterfactuals. To do so, we must specify the evolution of all fundamentals along this path.

We assume that local productivity evolves according to the damage function in (19), driven by the median temperature projections under the RCP 4.5 scenario (the solid line in Figure 6b). We further assume that all other fundamentals, such as trade and migration costs, remain constant at their 2014 levels. Given these assumptions, we solve for the deterministic equilibrium path using dynamic hat algebra (Caliendo et al., 2019). Intuitively, the 2014 allocation data pins down the initial levels of fundamentals, and the projected productivity changes then allow us to solve for the economy's entire future trajectory. In constructing

[28] Within the service sector, Cruz (2023) estimates the damage functions for construction, trade and transportation, finance, and government separately. We aggregate these coefficients based on their share in U.S. GDP to arrive at the coefficients for the broad service sector in our model. The estimates in Cruz (2023) are for value-added productivity. We scale these estimates by value-added shares to be consistent with our definition of z_{nj}^t as the TFP.

[29] In this application, we take the impact of climate change on fundamentals as given and abstract from the endogenous feedback from economic activity to carbon emissions in Integrated Assessment Models (e.g., Nordhaus, 1994). This simplification allows us to focus on the roles of uncertainty and heterogeneous beliefs. It is, however, possible to extend our model to incorporate endogenous temperature rise by adding a block on how carbon emission affects the temperature path.

these changes, we assume that the decisions of agents in 2014 were made under correct expectations about this median productivity path. We make this assumption primarily to create a clean benchmark that is comparable to the existing literature, which typically assumes perfect foresight. It is, however, conceptually straightforward to relax this assumption.[30]

4.2 Climate Change under Uncertainty and Heterogeneous Beliefs

We now analyze the economic impacts of climate change, focusing on two key departures from perfect foresight: aggregate uncertainty and heterogeneous beliefs. To isolate these effects, we first establish a perfectly anticipated climate change benchmark.

Perfectly anticipated climate change. Our benchmark scenario compares the deterministic path with climate change (driven by median temperature projections), to a no-change counterfactual where productivity remains at 2014 levels. The difference between these two paths captures the economic effect of perfectly anticipated climate change.

The red dotted lines in Figure 7 show the evolution of the population across locations in the United States with perfectly anticipated climate change. As the temperature rises, people move away from the Southwest. By 2100, the Southwest region accounts for only 8% of the U.S. population, down from 13% at the beginning of the period. In contrast, Western and Northeastern U.S. gain population. The blue solid lines in the figure show the evolution of the population across space under the 'no-change' scenario. We can see that if not for climate change, the population of the Southwest would increase and the population of the Northeast would decline.

Figure 8 quantifies the welfare effects of climate change. The green bars show the welfare change in each state under the perfectly anticipated scenario. The effects are highly polarized: northern states like Alaska and Maine experience large welfare gains, while southern states like Texas and Louisiana suffer losses of over 10%. These effects, both positive and negative, reflect the inverted-U relationship between temperature and productivity estimated in Table 1. On a population-weighted basis, the aggregate welfare loss for the U.S. is 1.6%, in the middle of the findings from other deterministic spatial models.[31]

[30] For example, if we had assumed agents in 2014 held different beliefs, we could first solve for the expected path consistent with those beliefs. Then, using Proposition 1, we could construct the 'true' median path as a deviation from that belief-driven path. In the same spirit, our framework can readily accommodate arbitrary time-varying paths for trade and migration costs.

[31] Rudik et al. (2022) estimates the damage function using the distribution of daily temperatures and infer a higher 3.1% welfare loss for the U.S. under RCP 4.5; on the other hand, Cruz and Rossi-Hansberg (2024) estimate a more modest loss than ours under RCP 6.0.

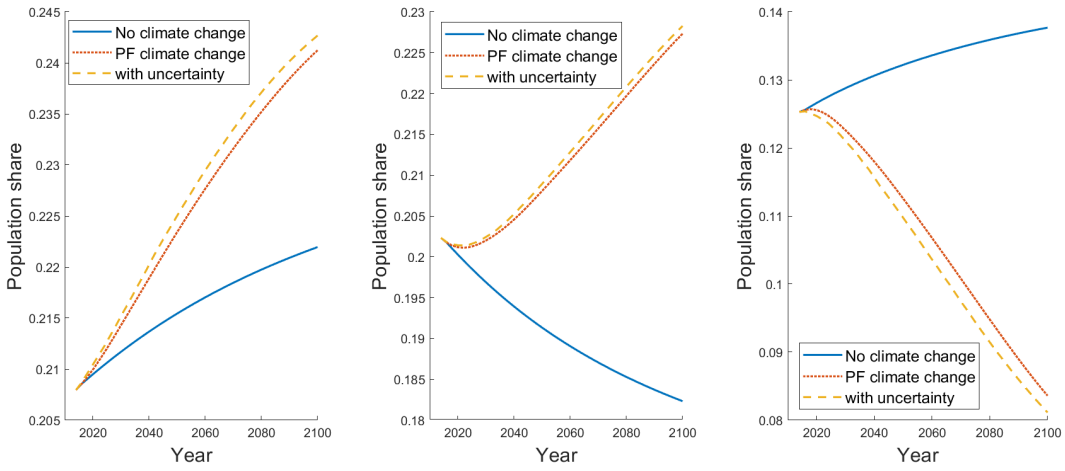

(a) West (b) Northeast (c) Southwest

Figure 7: Climate Change and Spatial Reallocation in the United States

Note: The figure shows each U.S. region's share of the population under different scenarios. 'No climate change' refers to the scenario in which climate and productivity stay the same as in 2014; 'PF climate change' refers to the scenario with perfectly anticipated climate change; and 'with uncertainty' refers to the scenario in which agents face uncertainty about future temperature rise and the effects of temperature on productivity.

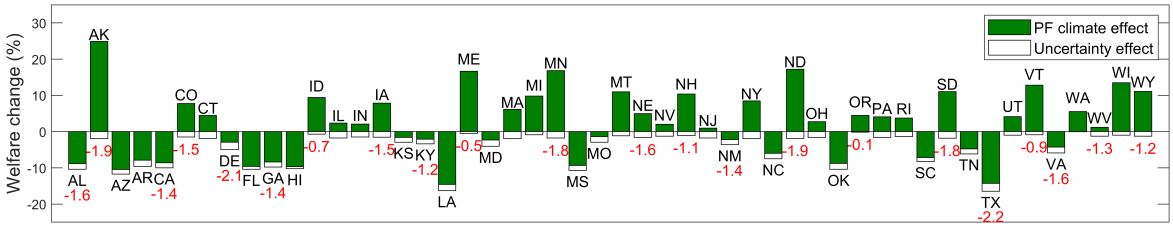

Figure 8: Welfare Effects of Climate Change

Note: The figure shows the welfare effect of climate change by U.S. state, measured in consumption-equivalent terms. The upper green bar represents the effect under perfect foresight, while the lower white bar captures the additional welfare loss due to uncertainty. Marked numbers indicate the magnitude of the uncertainty component.

The impact of uncertainty. We now introduce uncertainty. Agents in our model face uncertainty from two sources: the future path of temperature (confidence band in Figure 6b), and the damage function parameters themselves (the standard errors in Table 1).

To quantify the impact of such uncertainty, we apply the second-order approximation method from Proposition 3. For each period $t = 1, .., T$, we implement the second-order approximation around the deterministic path with climate change by simulating S future temperature paths. We then use the moments calculated from these S solutions to solve for agents' decision in period t, which shifts the economy into the next period.

To implement this procedure, we first specify the belief dynamics from which we generate temperature paths. We model agents' uncertainty today by adding mean-zero normal shocks around the median forecast for future temperatures. The variance of these shocks

$\sigma^1_{nt'}$, is calibrated to match the dispersion between the percentiles in Figure 6, with greater variance for more distant periods. Our assumption is essentially that agents today have similar information and beliefs as the scientific community on future temperature rise.

To solve the model forward, we must also specify how these beliefs will evolve. The challenge is that while the belief of the scientific community today can be measured, how such beliefs will be revised in the future as new data arrive is necessarily speculative. For example, if in the year 2030, the realization of temperature is higher than the median forecast for that year, will the future belief in 2030 be revised up or down? Will the perceived uncertainty in 2030 about future temperature be higher or lower? We assume that in any period t, agents' belief about the distribution for the period $t' \geq t$ is a normal distribution with mean $\mathcal{C}^t_{nt'}$ and variance $(\sigma^t_{nt'})^2$. We further assume that $\mathcal{C}^t_{nt'}$ follows a mean-reverting process that converges to the median forecast of temperature in t' by the scientific community today and that the uncertainty in period t' temperature perceived by workers standing at $t \leq t'$ depends only on $t' - t$. Formally,

$$\mathcal{C}^t_{nt'+1} - \bar{\mathcal{C}}_{nt'+1} = \rho(\mathcal{C}^t_{nt'} - \bar{\mathcal{C}}_{nt'}) \text{ and } \sigma^t_{nt'} = \sigma^1_{nt'-t+1}. \tag{20}$$

Here, $\sigma^1_{nt'}$ is calibrated before. The term $\bar{\mathcal{C}}_{nt'}$ is the median forecast for the period t' temperature from $t = 1$. This setup implies that if the realization of $\mathcal{C}_{nt}$ is above $\bar{\mathcal{C}}_{nt}$, then agents expect future temperatures to remain elevated but gradually return to $\bar{\mathcal{C}}_{nt'}$. We set $\rho = 0.9$, which implies a moderate convergence speed.

Second, to capture the spatial correlation in climate, we assume that all U.S. states experience the same percentile draw from their respective temperature distributions in any given period. For example, if Colorado experiences a 99th percentile shock, so does Utah.

Third, we evaluate the second-order solution along the median temperature path. While the stochastic equilibrium could, in principle, be solved along any temperature path, we choose the median path to ensure a clean comparison with our deterministic benchmark. This choice offers two key advantages. First, it means that as uncertainty vanishes, the stochastic solution converges exactly to the deterministic benchmark. Second, since the deterministic path is solved with a global method, this approach ensures that any difference we find is due to uncertainty, not due to effects from comparing second-order with linear approximations.

Fourth, to map temperature to productivity, we use the damage function from equation (19) to generate a productivity path for each simulation. We incorporate parameter uncertainty by drawing the coefficients, α^j_1 and α^j_2, from their estimated asymptotic distributions.

The orange dashed lines in Figure 7 show that uncertainty accelerates spatial reallocation. Migration out of the hot Southwest (Figure 7c) and into cooler regions like the West (Figure 7a) occurs faster than in the deterministic benchmark. By 2100, the Southwest's share of the U.S. population falls by 5.40 percentage points under perfect foresight, relative to a no-climate-change scenario. Uncertainty leads to an additional 0.25 percentage-point decline. This 'precautionary migration' is driven by risk aversion. The quadratic damage function means that in hotter regions, the marginal damage of rising temperature is larger. Therefore, the same variance in temperature uncertainty translates into a larger variance of future productivity. To avoid this uncertainty, agents move away from the South faster.

Incorporating uncertainty impacts the welfare loss from climate change. The white bars in Figure 8 show the additional welfare loss from uncertainty (in comparison to perfectly anticipated climate change). Because agents are risk-averse, this effect is negative in every state. The magnitude of this loss directly reflects the underlying productivity risk: it is largest in southern states like Texas (-2.2%), which face the highest productivity variance, and small in cooler states like Washington (-0.03%).

While this additional uncertainty effect—by construction second order—is modest for any individual state when compared to the effect from fully anticipated climate change, the fact that this effect is universally negative has important aggregate implications. Indeed, the large state-level effects from fully anticipated climate change partly offset each other, resulting in a 1.6% aggregate loss. In comparison, the additional aggregate loss from uncertainty is 1.48%. Incorporating uncertainty thus almost doubles the welfare costs of climate change.

In Appendix B.3.5, we report results from robustness exercises with lower risk aversion parameter, with smaller estimated effect of climate change on productivity, and with less persistent beliefs (lower ρ in equation (20)). Our finding that accounting for uncertainty substantially increases the welfare losses from climate change holds in all these cases.

Heterogeneous beliefs. We now consider another departure from perfect foresight—heterogeneous beliefs among agents. In Appendix Figure B.4, we show the spatial heterogeneity in beliefs about climate change. In particular, a larger share of the population are skeptical about climate change in the U.S. heartland than on the coasts.[32] We incorporate such heterogeneous beliefs by assuming that in 2014 each state has a share of population who do not believe in climate change and think the future temperature will be the same as

[32]This heterogeneity might be correlated with other observables, such as age or skill, which our model abstracts from. Our methodology works similarly if such group-specific heterogeneity is considered.

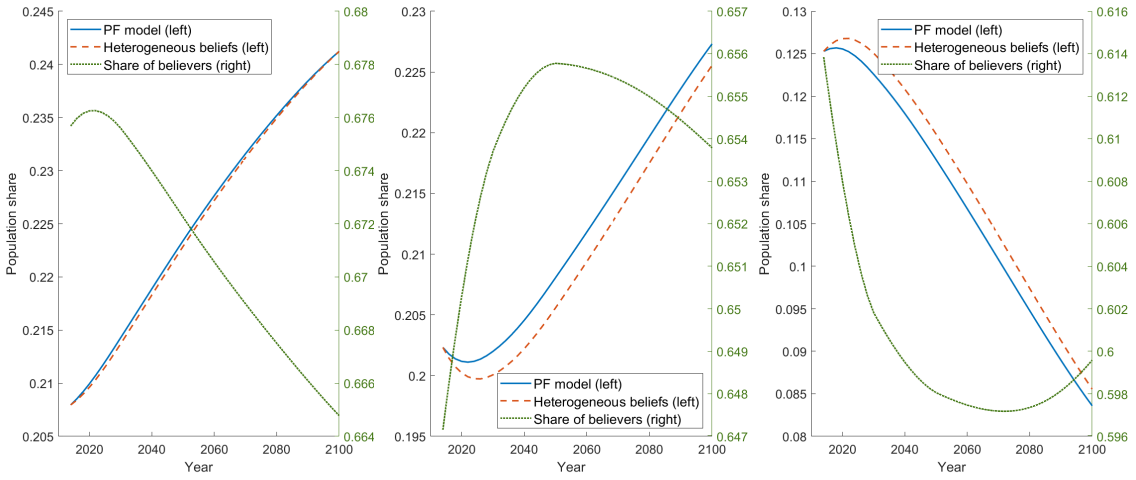

(a) West (b) Northeast (c) Southwest

Figure 9: Reallocation Across Space in the Presence of Non-believers

Note: The left axes show the evolution of the population shares under perfect foresight (solid blue lines) and under heterogenous beliefs (dashed lines). The right axes display the evolution of the share of believers across locations (dotted lines).

that in 2014. The remaining people believe in climate change and hold the belief that future productivity will evolve according to the median temperature path. The shares of both types are obtained from Figure B.4, and both types make location choices based on their beliefs about future outcomes, which depends on their beliefs about the other types' beliefs.

As discussed in Section 2.8, solving the model requires additional restrictions on higher-order beliefs. We adopt the first of the three assumptions on higher-order beliefs, discussed in Appendix A.4. Under this assumption, each type of agents, in making decisions, recognizes other types' beliefs, but believes that their own view will be vindicated and the other type will be convinced in the next period, when the new data arrive.[33] Under this assumption, we solve for the evolution of this economy in deviations from the perfect foresight baseline scenario (where everyone is a believer) following the procedure discussed in Appendix Proposition A.1.

Figure 9 plots the reallocation effect in the presence of climate skeptics in red dashed lines (the left axis). Compared to the baseline scenario in which everyone believes in climate change, here, the presence of climate skeptics results in a higher share of the population located in the South and a slower relocation of the population across space. Intuitively, as

[33]This assumption essentially truncates the third and high-order beliefs of agents. Recent studies show that in lab experiments of coordination games, most subjects engage in less than three levels of reasoning (Mauersberger and Nagel, 2018).

a fraction of the population does not believe the temperature will rise, these people tend to remain in, or move to, the South. Consistent with this intuition, the green dotted lines (the right axis) show that generally, the share of believers decreases in the South and increases in the Northeast.

In terms of welfare, compared with the baseline scenario, believers in the North and Midwest tend to be better off, whereas believers currently in the South tend to be worse off. The increase in the welfare of believers in northern locations stems from the more muted migration to the North, which reduces the labor market competition faced by workers there; the same force leads to a decrease in the welfare of believers in the South. The change in welfare is noticeable for some states; for example, the change is 1.6 percentage points in Maine and 0.5 percentage points in Wisconsin. In the aggregate, the welfare of believers increases by 0.02 percentage points. The welfare effects for individual states are displayed in Figure B.5 in the Appendix B.3.4.

5 Conclusion

The economic impact of shocks in dynamic spatial environments depends crucially on how agents form expectations about the future. While the perfect foresight assumption is common, it abstracts from the evolving, uncertain, and heterogeneous nature of beliefs. This paper develops a framework to analyze these environments.

Our work makes several contributions. We develop a recursive solution that solves an identification problem in ex-post analysis. It allows a researcher using observed data and belief information to tractably disentangle the role of agents' evolving beliefs from that of unobserved structural parameters like migration costs. This makes it feasible to conduct ex-post counterfactuals in settings where agents learn over time. We also develop a simulation-based second-order approximation method to quantify the effects of aggregate uncertainty. Our approach is flexible enough to handle time-varying, history-dependent stochastic processes and can be extended to incorporate features like heterogeneous beliefs. Finally, we provide a framework for conducting theory-consistent welfare analysis in these settings, showing how to measure welfare both ex-ante and ex-post.

To showcase the methodology, we conduct two applications: an ex-post study of the impact of China's productivity growth on U.S. employment, and an ex-ante study of the spatial reallocation in response to climate change. Both quantitative applications demonstrate the importance of incorporating the departures in agents' beliefs from perfect foresight.

References

Alessandria, George A, Shafaat Y Khan, Armen Khederlarian, Kim J Ruhl, and Joseph B Steinberg, "Trade-Policy Dynamics: Evidence from 60 Years of US–China Trade," *Journal of Political Economy*, 2025, *133* (3), 713–749.

Armantier, Olivier, Giorgio Topa, Wilbert Van der Klaauw, and Basit Zafar, "An Overview of the Survey of Consumer Expectations," *Federal Reserve Bank of New York Economic Policy Review*, 2017, *23* (2), 51–72.

Artuç, Erhan, Shubham Chaudhuri, and John McLaren, "Trade Shocks and Labor Adjustment: A Structural Empirical Approach," *American Economic Review*, 2010, *100* (3), 1008–45.

Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub, "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models," *Econometrica*, 2021, *89* (5), 2375–2408.

Autor, David H, David Dorn, and Gordon H Hanson, "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 2013, *103* (6), 2121–2168.

Balboni, Clare, "In Harm's Way? Infrastructure Investments and the Persistence of Coastal Cities," *American Economic Review*, 2025, *115* (1), 77–116.

Bhandari, Anmol, Thomas Bourany, David Evans, and Mikhail Golosov, "A Perturbational Approach for Approximating Heterogeneous Agent Models," *Working Paper*, 2023.

Bilal, Adrien, "Solving Heterogeneous Agent Models with the Master Equation," *NBER Working Paper No. 31103*, 2023.

Bombardini, Matilde, Bingjing Li, and Francesco Trebbi, "Did US Politicians Expect the China Shock?," *American Economic Review*, 2023, *113* (1), 174–209.

Bui, Ha, Zhen Huo, Andrei A Levchenko, and Nitya Pandalai-Nayar, "Information Frictions and News Media in Global Value Chains," *NBER Working Paper No. 30033*, 2022.

Caliendo, Lorenzo and Fernando Parro, "Estimates of the Trade and Welfare Effects of NAFTA," *The Review of Economic Studies*, 2015, *82* (1), 1–44.

_ , Maximiliano Dvorkin, and Fernando Parro, "Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock," *Econometrica*, 2019, *87* (3), 741–835.

Cogley, Timothy and Thomas J Sargent, "Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making," *International Economic Review*, 2008, *49* (1), 185–221.

Coibion, Olivier, Yuriy Gorodnichenko, and Saten Kumar, "How Do Firms Form Their Expectations? New Survey Evidence," *American Economic Review*, 2018, *108* (9), 2671–2713.

Cruz, José-Luis, "Global Warming and Labor Market Reallocation," *Unpublished Manuscript*, 2023.

Cruz, José-Luis and Esteban Rossi-Hansberg, "The Economic Geography of Global Warming," *Review of Economic Studies*, 2024, *91* (2), 899–939.

Dekle, Robert, Jonathan Eaton, and Samuel Kortum, "Unbalanced Trade," *American Economic Review Papers and Proceedings*, 2007, *97* (2), 351–355.

Desmet, Klaus and Fernando Parro, "Spatial Dynamics," Working Paper 33443, National Bureau of Economic Research 2025.

Dickstein, Michael J and Eduardo Morales, "What Do Exporters Know?," *The Quarterly Journal of Economics*, 2018, *133* (4), 1753–1801.

Dix-Carneiro, Rafael, "Trade Liberalization and Labor Market Dynamics," *Econometrica*, 2014, *82* (3), 825–885.

_ , João Paulo Pessoa, Ricardo M Reyes-Heroles, and Sharon Traiberman, "Globalization, Trade Imbalances and Labor Market Adjustment," *Quarterly Journal of Economics*, 2023, *138* (2), 1109–1171.

Eaton, Jonathan and Samuel Kortum, "Technology, Geography, and Trade," *Econometrica*, 2002, *70* (5), 1741–1779.

Egger, Peter, Katharina Erhardt, and Davide Suverato, "How Aggregate Uncertainty Shapes the Spatial Economy," *Unpublished Manuscript*, 2024.

Fan, Jingting and Wenlan Luo, "Global Production Networks with Global Uncertainty," *Available at SSRN 5282371*, 2025.

Galle, Simon, Andrés Rodríguez-Clare, and Moises Yi, "Slicing the Pie: Quantifying the Aggregate and Distributional Effects of Trade," *The Review of Economic Studies*, 2023, *90* (1), 331–375.

Giannone, Elisa, Qi Li, Nuno Paixao, and Xinle Pang, "Unpacking Moving: A Quantitative Spatial Equilibrium Model with Wealth," *Bank of Canada Staff Working Paper*, 2023.

Greenwood, Robin and Andrei Shleifer, "Expectations of Returns and Expected Returns," *The Review of Financial Studies*, 2014, *27* (3), 714–746.

Handley, Kyle and Nuno Limão, "Policy Uncertainty, Trade, and Welfare: Theory and Evidence for China and the United States," *American Economic Review*, 2017, *107* (9), 2731–2783.

Huo, Zhen and Naoki Takayama, "Rational Expectations Models with Higher-Order Beliefs," *Review of Economic Studies*, 2024, p. 3138–3173.

Kim, Jinill, Sunghyun Kim, Ernst Schaumburg, and Christopher A Sims, "Calculating and Using Second-Order Accurate Solutions of Discrete Time Dynamic Equilibrium Models," *Journal of Economic Dynamics and Control*, 2008, *32* (11), 3397–3414.

Kleinman, Benny, Ernest Liu, and Stephen J Redding, "Dynamic Spatial General Equilibrium," *Econometrica*, 2023, *91* (2), 385–424.

_ , _ , and _ , "International Trade in an Uncertain World," *Working Paper*, 2025.

Kozlowski, Julian, Laura Veldkamp, and Venky Venkateswaran, "The Tail that Wags the Economy: Beliefs and Persistent Stagnation," *Journal of Political Economy*, 2020, *128* (8), 2839–2879.

Kreps, David M, "Anticipated Utility and Dynamic Choice," *Econometric Society Monographs*, 1998, *29*, 242–274.

Krusell, Per and Anthony Smith, "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, October 1998, *106* (5), 867–896.

Malmendier, Ulrike and Stefan Nagel, "Learning from Inflation Experiences," *The Quarterly Journal of Economics*, 2016, *131* (1), 53–87.

Mauersberger, Felix and Rosemarie Nagel, "Levels of Reasoning in Keynesian Beauty Contests: A Generative Framework," in "Handbook of Computational Economics," Vol. 4, Elsevier, 2018, pp. 541–634.

Mennuni, Alessandro, Juan F Rubio-Ramirez, and Serhiy Stepanchuk, "Dynamic Perturbation," *Review of Economic Studies*, 2024, p. 1157–1192.

Nordhaus, William, *Managing the Global Commons: the Economics of Climate Change*, The MIT Press, 1994.

Pang, Xinle and Sun Pin, "Moving into Risky Floodplains: The Spatial Implications of Federal Disaster Relief Policies," *Unpublished Manuscript*, 2022.

Pierce, Justin R and Peter K Schott, "The Surprisingly Swift Decline of US Manufacturing Employment," *American Economic Review*, 2016, *106* (7), 1632–1662.

Porcher, Charly, "Migration with Costly Information," *Working Paper*, 2022.

_ , Eduardo Morales, and Thomas Fujiwara, "Measuring Information Frictions in Migration: A Revealed-Preference Approach," *Working Paper*, 2025.

Rodríguez-Clare, Andrés, Mauricio Ulate, and Jose P. Vásquez, "Trade with Nominal Rigidities: Understanding the Unemployment and Welfare Effects of the China Shock," *Journal of Political Economy, forthcoming*, 2025.

Rudik, Ivan, Gary Lyn, Weiliang Tan, and Ariel Ortiz-Bobea, "The Economic Effects of Climate Change in Dynamic Spatial Equilibrium," *Working Paper*, 2022.

Schmitt-Grohé, Stephanie and Martın Uribe, "Solving Dynamic General Equilibrium Models Using a Second-Order Approximation to the Policy Function," *Journal of Economic Dynamics and Control*, 2004, *28* (4), 755–775.

Townsend, Robert M, "Forecasting the Forecasts of Others," *Journal of Political Economy*, 1983, *91* (4), 546–588.

Uhlig, Harald, "A Toolkit for Analyzing Nonlinear Dynamic Stochastic Models Easily," *Federal Reserve Bank of Minneapolis Discussion Paper*, 1995.

Online Appendix: Learning and Expectations in Dynamic Spatial Economies

Jingting Fan
Cheung Kong Graduate School of Business

Sungwan Hong
University of Pittsburgh

Fernando Parro
University of Rochester and NBER

Appendix A Theory

This appendix presents the proofs of the propositions and additional derivations. To simplify the notation, we denote a conditional expectation by $\mathbb{E}_t x_{nt'} \equiv \mathbb{E}[x_{nt'}(z^{t'})|z^t]$. Throughout, we define $\hat{v}_{it+1}(z^t) \equiv v_{it+1}(z^t) - \bar{v}_{it+1}$ and $\hat{v}_{it}(z^t) \equiv v_{it}(z^t) - \bar{v}_{it}$; for all other exogenous or endogenous variables x, we define $\hat{x}_t(z^t) \equiv \ln x_t(z^t) - \ln \bar{x}_t$. The proofs and additional derivations in this appendix are based on the one-sector version of the model; the step-by-step algebra for the multi-sector version of the model used in the quantification sections is relegated to an Online Supplement.

To demonstrate that the methodology can accommodate other fundamentals besides productivity being stochastic, in all proofs and derivations provided in this appendix, we allow trade costs κ_{nit} to also be stochastic. Stochastic migration costs can be incorporated analogously. In the quantitative applications, we maintain that trade and migration costs are perfectly anticipated.

A.1 Proof of Proposition 1

We prove the proposition by deriving a first-order approximation of the model's equilibrium conditions (equations (3)-(8)). We illustrate the main steps using the location value (equation (3)) as an example. First, start with $\hat{v}_{it+1}(z^t)$, the difference between the continuation value and its deterministic counterpart:

$$\begin{aligned}
\hat{v}_{it+1}(z^t) \equiv & \int_\Theta [U(c_{it+1}(z^{t+1})) + \nu \ln(\sum_{n=1}^N \exp(\beta v_{nt+2}(z^{t+1}) - m_{int+1})^{1/\nu})] f(z_{t+1}|z^t) dz_{t+1} \\
& - [U(\bar{c}_{it+1}) + \nu \ln(\sum_{n=1}^N \exp(\beta \bar{v}_{nt+2} - m_{int+1})^{1/\nu})] \\
\equiv & \int_\Theta [F_i(c_{it+1}(z^{t+1}), v_{t+2}(z^{t+1})) - F_i(\bar{c}_{it+1}, \bar{v}_{t+2})] f(z_{t+1}|z^t) dz_{t+1} \\
= & \int_\Theta [\frac{\partial F_i}{\partial \bar{c}_{it+1}} \hat{c}_{it+1}(z^{t+1}) + \frac{\partial F_i}{\partial \bar{v}_{t+2}} \hat{v}_{t+2}(z^{t+1}) + o(\hat{v}_{t+2}(z^{t+1}), \hat{c}_{it+1}(z^{t+1}))] f(z_{t+1}|z^t) dz_{t+1},
\end{aligned}$$

where we have defined the terms inside the bracket in the first line as a function F_i. Thus, $\hat{v}_{it+1}(z^t)$ can be seen as the difference between F_i evaluated at the stochastic path versus the deterministic path, integrated over possible realizations in $t+1$. The last line of the equation is a first-order Taylor approximation of $\left[F_i\left(c_{it+1}(z^{t+1}), v_{t+2}(z^{t+1})\right) - F_i\left(\bar{c}_{it+1}, \bar{v}_{t+2}\right)\right]$ at

$(\bar{c}_{it+1}, \bar{v}_{t+2})$, with a second-order approximation error $o(\hat{v}_{t+2}(z^{t+1}), \hat{c}_{it+1}(z^{t+1}))$.[1] Integrating $\hat{v}_{it+1}(z^t)$ over possible paths of fundamental productivities up to t from the perspective of an agent making decisions at $t = 1$, we obtain

$$\begin{aligned}
&\int_{\Theta^t} \hat{v}_{it+1}(z^t) f(z^t|z_1) dz^t \equiv \mathbb{E}_1 \hat{v}_{it+1} \approx \frac{\partial F_i}{\partial \bar{c}_{it+1}} \mathbb{E}_1 \hat{c}_{it+1} + \frac{\partial F_i}{\partial \bar{v}_{t+2}} \mathbb{E}_1 \hat{v}_{t+2} \\
&\implies \mathbb{E}_1 \hat{v}_{it+1} \approx \frac{\partial U(\bar{c}_{it+1})}{\partial \ln \bar{c}_{it+1}} \mathbb{E}_1 \hat{c}_{it+1} + \beta \sum_n \bar{\mu}_{int+1} \mathbb{E}_1 \hat{v}_{nt+2}. \qquad \text{(A.1)}
\end{aligned}$$

The second line evaluates $\frac{\partial F_i}{\partial \bar{c}_{it+1}}$ and $\frac{\partial F_i}{\partial \bar{v}_{t+2}}$ (see Online Supplement C.2 for step-by-step algebra). In particular, $\frac{\partial U(\bar{c}_{it})}{\partial \ln \bar{c}_{it+1}}$ is the elasticity of flow utility to real consumption; $\bar{\mu}_{int+1}$ are the migration shares on the deterministic approximation path. Equation (A.1) therefore expresses the expectation (in eyes of agents in $t = 1$) of the deviation from perfect foresight in period $t + 1$ location values as a linear function of $\mathbb{E}_1 \hat{c}_{it+1}$ and $\{\mathbb{E}_1 \hat{v}_{nt+2}\}_{n=1}^N$, with real consumption and migration shares on the deterministic path as their coefficients.

We approximate analogously the expected values of other endogenous variables around their counterparts on the deterministic sequential path. In each of these approximations, we can derive the partial derivatives (that are akin to $\frac{\partial F_i}{\partial \bar{c}_{it+1}}$ and $\frac{\partial F_i}{\partial \bar{v}_{t+2}}$) as closed-form functions of the outcomes on the deterministic path with perfect foresight, obtaining the system of linear equations for $t' \geq t$. The derivations for the rest of the equation system are analogous; since they are straightforward, we delegate the step-by-step algebra to the Online Supplement C.2.

Algorithm. To solve the system of linear equations, we iterate on the guesses of the path of expected location values according to agents' belief in period t ($\mathbb{E}_t \hat{v}_{nt'}$). Given a guess of these values and the initial labor allocation, we use equations (10) and (11) to obtain the expected path of the population distribution. With this path, we solve for the static trade outcomes in each t' using equations (12)-(14). We construct the path of expected real wages using the solution and then update the guess on the path of expected location values using (9). We iterate on this process until convergence, which delivers agents' actual migration decisions at time t and therefore moves the economy to the next period through the law of motion for labor. We then repeat the procedure to solve for the actual allocations at each time $t + 1$, $t + 2$,..., to obtain the evolution of the economy. See Online Supplement D.5 for additional details on numerical implementation of the algorithm.

This process requires solving the static trade equilibrium numerous times. We derive

[1] In $\frac{\partial F_i}{\partial \bar{x}}$, we use $\partial \bar{x}$ to denote level derivatives with respect to x when x is the value of locations, and to denote log derivatives with respect to x when x is other variables. In both cases, the derivatives are evaluated at the deterministic sequential equilibrium, hence the bar in $\partial \bar{x}$.

step-by-step in Online Supplement D.6 the analytical expressions of outcomes in the static trade equilibrium ($\mathbb{E}_t\hat{\lambda}_{nit'}, \mathbb{E}_t\hat{P}_{it'}, \mathbb{E}_t\hat{w}_{it'}$) for an extended model with multiple sectors and intermediate inputs, so solving the static trade equilibrium is a simple matrix inversion.

A.2 Proof of Proposition 2

In this subsection, we first explain the recursive solution presented in Section 2.5 for recovering belief and perfect foresight paths from the observed allocation. We then describe how we use the expected paths and the actual allocation to measure welfare when agents do not have perfect foresight.

Recovering belief and perfect foresight paths. We use variables with a bar $\bar{}$ to denote the approximation points, which we will construct sequentially. We denote the *actual* outcomes in period t by variables with a tilde $\tilde{}$, i.e., $\tilde{v}_{nt}$, $\tilde{w}_{nt}$, $\tilde{P}_{nt}$, $\tilde{\lambda}_{nit}$, $\tilde{\mu}_{nit}$, $\tilde{l}_{nt}$.

i. Starting from period T, we solve for the expected outcomes of the trade equilibrium in period T according to the beliefs of agents in period $T-1$, namely, $\mathbb{E}_{T-1}w_{iT}$, $\mathbb{E}_{T-1}P_{iT}$, and $\mathbb{E}_{T-1}\lambda_{inT}$ for all i, n. We do so by approximating around the actual outcomes in period T. Define $\mathbb{E}_{T-1}\hat{x}_T \equiv \mathbb{E}_{T-1}\ln(x_T) - \ln(\tilde{x}_T)$ for $x \in \{w, P, \lambda, l\}$. These hat variables satisfy the following system of linear equations:

$$\begin{aligned}
\mathbb{E}_{T-1}\hat{\lambda}_{niT} &= -\theta\Big(\mathbb{E}_{T-1}\hat{w}_{iT} + \mathbb{E}_{T-1}\hat{\kappa}_{niT} - \mathbb{E}_{T-1}\hat{P}_{nT}\Big) + \mathbb{E}_{T-1}\hat{z}_T \\
\mathbb{E}_{T-1}\hat{P}_{nT} &= \sum_{i=1}^{N}\bar{\lambda}_{niT}\Big(\mathbb{E}_{T-1}\hat{w}_{iT} + \mathbb{E}_{T-1}\hat{\kappa}_{niT} - \frac{1}{\theta}\mathbb{E}_{T-1}\hat{z}_T\Big) \qquad \text{(A.2)} \\
\mathbb{E}_{T-1}\hat{w}_{nT} + \mathbb{E}_{T-1}\hat{l}_{nT} &= \sum_{i=1}^{N}\frac{\bar{\lambda}_{inT}\bar{w}_{iT}\bar{l}_{iT}}{\bar{w}_{nT}\bar{l}_{nT}}\Big(\mathbb{E}_{T-1}\hat{\lambda}_{inT} + \mathbb{E}_{T-1}\hat{w}_{iT} + \mathbb{E}_{T-1}\hat{l}_{iT}\Big),
\end{aligned}$$

where the approximation points $(\bar{\lambda}_{inT}, \bar{w}_{nT}, \bar{P}_{nT}, \bar{l}_{nT})$ are defined to be the actual outcomes in T, namely $(\tilde{\lambda}_{inT}, \tilde{w}_{nT}, \tilde{P}_{nT}, \tilde{l}_{nT})$. Moreover, since the actual labor allocation in T is determined by the migration decision at $T-1$, we can plug in $\mathbb{E}_{T-1}\hat{l}_{iT} = 0$. The input for this system of equations is the actual outcomes in T and the deviations in agents' belief in period $T-1$ about the productivity in T from the actual productivity, namely, $\mathbb{E}_{T-1}\hat{z}_T$ (and in the case of stochastic trade costs, $\mathbb{E}_{T-1}\hat{\kappa}_{inT}$). The output of this system of equations is deviations such as $\mathbb{E}_{T-1}\hat{\lambda}_{niT}$, with which we can recover the level outcomes using the following relationship: $\mathbb{E}_{T-1}\ln(x_T) = \mathbb{E}_{T-1}\hat{x}_T + \ln(\tilde{x}_T)$ for $x \in \{w, P, \lambda, l\}$.

ii. Append the output from the first step with $\{\tilde{\mu}_{niT-1}, \tilde{w}_{nT-1}, \tilde{P}_{iT-1}, \tilde{\lambda}_{iT-1}, \tilde{l}_{nT-1}\}$, namely, agents' *actual* migration decision and the actual outcomes of the static trade equilibrium in period $T-1$. As $\{\tilde{\mu}_{niT-1}, \tilde{w}_{nT-1}, \tilde{P}_{iT-1}, \tilde{\lambda}_{iT-1}\}$ are all determined in period $T-1$, when the agents have observed the realization of productivity in $T-1$, together with the output from (i), they constitute the solution to the agents' problem in period $T-1$, as defined in Proposition 1.
Steps (i) and (ii) delivers $\{\mathbb{E}_{T-1}\lambda_{nit}, \mathbb{E}_{T-1}P_{nt}, \mathbb{E}_{T-1}w_{nt}, \mathbb{E}_{T-1}l_{nt}, \mathbb{E}_{T-1}\mu_{nit}\}_{i=1,n=1,t=T-1}^{N,N,T}$, namely, the outcomes in $\{T-1, T\}$ according to agents' belief in period $T-1$.

iii. Derive the expected outcomes in periods $\{T-1, T\}$ according to agents' belief in period $T-2$ by approximating around the solution to agents' problem in period $T-1$ obtained previously. Doing so delivers a system of linear equations (see below) with the input being (a) the outputs from steps (i) and (ii), and (b) the deviations in agents' beliefs in period $T-2$ about the productivity in $\{T-1, T\}$ from their beliefs in period $T-1$ about it, namely, $\mathbb{E}_{T-2}\ln(z_t) - \mathbb{E}_{T-1}\ln(z_t)$ for $t \in \{T-1, T\}$.
The output of the equations is $\{\mathbb{E}_{T-2}\hat{\lambda}_{nit}, \mathbb{E}_{T-2}\hat{P}_{nt}, \mathbb{E}_{T-2}\hat{w}_{nt}, \mathbb{E}_{T-2}\hat{l}_{nt}, \mathbb{E}_{T-2}\hat{\mu}_{nit}\}_{i=1,n=1,t=T-1}^{N,N,T}$, with which we can recover the corresponding level variables. Here, the expectation operators stand for the difference in the expected value of an outcome between expectations formed in $T-2$ and the expectation formed in $T-1$. For example, $\mathbb{E}_{T-2}\hat{v}_t \equiv \mathbb{E}_{T-2}v_t - \mathbb{E}_{T-1}\tilde{v}_t$ for $t \in \{T-1, T\}$.
The system of equations for $t \in \{T-1, T\}$ is given by

$$
\begin{aligned}
\mathbb{E}_{T-2}\hat{\lambda}_{nit} &= -\theta\Big(\mathbb{E}_{T-2}\hat{w}_{it} + \mathbb{E}_{T-2}\hat{\kappa}_{nit} - \mathbb{E}_{T-2}\hat{P}_{nt}\Big) + \mathbb{E}_{T-2}\hat{z}_t \\
\mathbb{E}_{T-2}\hat{P}_{nt} &= \sum_{i=1}^{N} \bar{\lambda}_{nit}\Big(\mathbb{E}\hat{w}_{it} + \mathbb{E}_{T-2}\hat{\kappa}_{nit} - \frac{1}{\theta}\mathbb{E}_{T-2}\hat{z}_t\Big) \\
\mathbb{E}_{T-2}\hat{w}_{nt} + \mathbb{E}_{T-2}\hat{l}_{nt} &= \sum_{i=1}^{N} \frac{\bar{\lambda}_{int}\bar{w}_{it}\bar{l}_{it}}{\bar{w}_{nt}\bar{l}_{nt}}\Big(\mathbb{E}_{T-2}\hat{\lambda}_{int} + \mathbb{E}_{T-2}\hat{w}_{it} + \mathbb{E}_{T-2}\hat{l}_{it}\Big) \\
\mathbb{E}_{T-2}\hat{v}_{nt} &= \frac{\partial U(\bar{w}_{nt}, \bar{P}_{nt})}{\partial \ln(\bar{w}_{nt})}\mathbb{E}_{T-2}\hat{w}_{nt} + \frac{\partial U(\bar{w}_{nt}, \bar{P}_{nt})}{\partial \ln(\bar{P}_{nt})}\mathbb{E}_{T-2}\hat{P}_{nt} + \beta\sum_{i=1}^{N}\bar{\mu}_{nit}\mathbb{E}_{T-2}\hat{v}_{it+1} \quad \text{(A.3)} \\
\mathbb{E}_{T-2}\hat{\mu}_{nit} &= \frac{\beta}{\nu}\sum_{m=1}^{N}\Big(\mathbb{1}(m=i) - \bar{\mu}_{nmt}\Big)\mathbb{E}_{T-2}\hat{v}_{mt+1} \\
\mathbb{E}_{T-2}\hat{l}_{nt+1} &= \sum_{i=1}^{N}\frac{\bar{\mu}_{int}\bar{l}_{it}}{\bar{l}_{nt+1}}\Big(\mathbb{E}_{T-2}\hat{\mu}_{int} + \mathbb{E}_{T-2}\hat{l}_{it}\Big).
\end{aligned}
$$

The approximation points are $\{\mathbb{E}_{T-1}\lambda_{nit}, \mathbb{E}_{T-1}P_{nt}, \mathbb{E}_{T-1}w_{nt}, \mathbb{E}_{T-1}l_{nt}, \mathbb{E}_{T-1}\mu_{nit}\}_{i=1,n=1,t=T-1}^{N,N,T}$.

For example, $\bar{\lambda}_{int} \equiv \mathbb{E}_{T-1}\lambda_{nit}$ and $\bar{\mu}_{int} \equiv \mathbb{E}_{T-1}\mu_{nit}$. Using the inputs described above and the fact that $\mathbb{E}_{T-2}\hat{l}_{iT-1}$=0, we can solve the system of the equations to obtain $\{\mathbb{E}_{T-2}\hat{\lambda}_{nit}, \mathbb{E}_{T-2}\hat{P}_{nt}, \mathbb{E}_{T-2}\hat{w}_{nt}, \mathbb{E}_{T-2}\hat{l}_{nt}, \mathbb{E}_{T-2}\hat{\mu}_{nit}, \mathbb{E}_{T-2}\hat{v}_t\}_{i=1,n=1,t=T-1}^{N,N,T}$.

iv. Append the output from (iii) by $\{\tilde{\mu}_{niT-2}, \tilde{w}_{nT-2}, \tilde{P}_{iT-2}, \tilde{\lambda}_{iT-2}\}$, namely, agents' *actual* migration decision and the actual outcomes of the static trade equilibrium in period $T-2$. Together with the output from the previous step, this constitutes the solution to the agents' problem in period $T-2$.

v. Note that steps (iii) and (iv) take as input the beliefs of agents in period $T-1$ about outcomes in $t \in \{T-1, T\}$ and produce as output the belief of agents in period $T-2$ about outcomes in $t \in \{T-2, T-1, T\}$. Iterating on steps (iii) and (iv) backward recursively until the first period delivers the expected paths in all periods.

With the expected paths at hand, we can approximate around each of these paths to obtain the perfect foresight path for agents in the corresponding periods. Both the expected and perfect foresight paths can serve as an approximation path to recover counterfactual outcomes.

Extension for when the data end at $T' < T$. In the above discussion we have assumed that the researcher observes the actual allocation until the end of model T. Suppose the researcher has only observed the outcome up until $T' < T$. We can proceed by first constructing the expected path of agents in period T' using a time-difference version of the model, which can be solved using the time-difference algorithm described in Caliendo et al. (2019) with first-order accuracy.[2] Doing so requires taking a stand on the period-by-period changes in the fundamentals between T' and T according to agents' belief in T', which is inevitable as no data are observed beyond T'.[3] With the expected path in T' at hand, iterating on steps (iii) and (iv) delivers the expected paths from all periods.

A.3 Proof of Proposition 3

We present the system of equations for second-order solution and the solution algorithm. Step-by-step algebra for the terms is delegated to the Online Supplement C.3.

[2] If beliefs are deterministic, then the solution is exact.

[3] Intuitively, the observed data contain information on the level of fundamentals in T'. This information, coupled with the assumption on future changes in these fundamentals, enables the researcher to solve for the evolution according to agents' beliefs.

A second-order approximation of equations (3)-(8) that characterizes the decisions and future expected outcomes for agents in period t is [4]

$$
\begin{aligned}
\mathbb{E}_t\hat{v}_{nt'} &= \frac{\partial U(\bar{w}_{nt'},\bar{P}_{nt'})}{\partial \ln \bar{w}_{nt'}}\mathbb{E}_t\hat{w}_{nt'} + \frac{\partial U(\bar{w}_{nt'},\bar{P}_{nt'})}{\partial \ln \bar{P}_{nt'}}\mathbb{E}_t\hat{P}_{nt'} + \frac{1}{2}\frac{\partial^2 U}{\partial \ln \bar{w}_{nt'}^2}\mathbb{E}_t\hat{w}_{nt'}^2 \\
&+\frac{1}{2}\frac{\partial^2 U}{\partial \ln \bar{P}_{nt'}^2}\mathbb{E}_t\hat{P}_{nt'}^2 + \frac{\partial^2 U}{\partial \ln \bar{w}_{nt'}\partial \ln \bar{P}_{nt'}}\mathbb{E}_t\hat{w}_{nt'}\hat{P}_{nt'} + \sum_m \beta\bar{\mu}_{nmt'}\mathbb{E}_t\hat{v}_{mt'+1} \\
&+\frac{1}{2}\sum_{o,m}\frac{\beta^2}{\nu}\bar{\mu}_{not'}\left[\mathbf{1}(m=o)-\bar{\mu}_{nmt'}\right] \\
\mathbb{E}_t\hat{\mu}_{nit'} &= \frac{\beta}{\nu}\sum_m\left[\mathbf{1}(m=i)-\bar{\mu}_{nmt'}\right]\mathbb{E}_t\hat{v}_{mt'+1} + \frac{1}{2}\left(\frac{\beta}{\nu}\right)^2\sum_{m,o}\Big\{\mathbf{1}(m=i)\bar{\mu}_{nit'}\left[\mathbf{1}(o=i)-\bar{\mu}_{not'}\right] \\
&-\bar{\mu}_{nit'}\bar{\mu}_{nmt'}\left[\mathbf{1}(o=i)-2\bar{\mu}_{not'}+\mathbf{1}(o=m)\right]\Big\}\mathbb{E}_t\hat{v}_{mt'+1}\hat{v}_{ot'+1}, \\
\mathbb{E}_t\hat{l}_{nt'+1} &= \sum_{i,m}\frac{\beta}{\nu}\bar{\psi}_{int'+1}\left[\mathbf{1}(m=n)-\bar{\mu}_{imt'}\right]\mathbb{E}_t\hat{v}_{mt'+1} + \sum_m\bar{\psi}_{mnt'+1}\mathbb{E}_t\hat{l}_{mt'} \qquad \text{(A.4)} \\
&+\frac{1}{2}\sum_{m,o}\left(\frac{\beta}{\nu}\right)^2\Big[-\sum_i\bar{\mu}_{imt'}\bar{\psi}_{int'+1}\left[\mathbf{1}(o=n)+\mathbf{1}(o=m)-2\bar{\mu}_{iot'}\right] \\
&+\mathbf{1}(m=n)\sum_i\bar{\psi}_{int'+1}\left[\mathbf{1}(o=n)-\bar{\mu}_{iot'}\right] \\
&-\Big(\sum_i\bar{\psi}_{int'+1}\left[\mathbf{1}(m=n)-\bar{\mu}_{imt'}\right]\Big)\Big(\sum_i\bar{\psi}_{int'+1}\left[\mathbf{1}(o=n)-\bar{\mu}_{iot'}\right]\Big)\Big]\mathbb{E}_t\hat{v}_{mt'+1}\hat{v}_{ot'+1} \\
&+\frac{1}{2}\sum_{m,o}\left[\mathbf{1}(m=o)\bar{\psi}_{ont'+1}-\bar{\psi}_{mnt'+1}\bar{\psi}_{ont'+1}\right]\mathbb{E}_t\hat{l}_{mt'}\hat{l}_{ot'} \\
&+\sum_{m,o}\bar{\psi}_{ont'+1}\frac{\beta}{\nu}\Big[\left[\mathbf{1}(o=n)-\bar{\mu}_{omt'}\right]-\sum_i\bar{\psi}_{int'+1}\left[\mathbf{1}(m=n)-\bar{\mu}_{imt'}\right]\Big]\mathbb{E}_t\hat{v}_{mt'+1}\hat{l}_{ot'}, \\
\mathbb{E}_t\hat{\lambda}_{nit'} &= -\theta\left(\mathbb{E}_t\hat{w}_{it'}+\mathbb{E}_t\hat{\kappa}_{nit'}-\mathbb{E}_t\hat{P}_{nt'}\right)+\mathbb{E}_t\hat{z}_{it'}, \\
\mathbb{E}_t\hat{P}_{nt'} &= \sum_i\bar{\lambda}_{nit'}\mathbb{E}_t\hat{w}_{it'}+\sum_i\bar{\lambda}_{nit'}\mathbb{E}_t\hat{\kappa}_{nit'}-\frac{1}{\theta}\sum_i\bar{\lambda}_{nit'}\mathbb{E}_t\hat{z}_{it'} \\
&+\frac{1}{2}\sum_{i,m}\theta\left[-\mathbf{1}(i=m)\bar{\lambda}_{nit'}+\bar{\lambda}_{nit'}\bar{\lambda}_{nmt'}\right]\left(\mathbb{E}_t\hat{w}_{it'}\hat{w}_{mt'}+\mathbb{E}_t\hat{\kappa}_{nit'}\hat{\kappa}_{nmt'}+\theta^{-2}\mathbb{E}_t\hat{z}_{it'}\hat{z}_{mt'}\right) \\
&+\sum_{i,m}\theta\left[-\mathbf{1}(i=m)\bar{\lambda}_{nit'}+\bar{\lambda}_{nit'}\bar{\lambda}_{nmt'}\right]\left(\mathbb{E}_t\hat{w}_{it'}\hat{\kappa}_{nmt'}-\theta^{-1}\mathbb{E}_t\hat{w}_{it'}\hat{z}_{mt'}-\theta^{-1}\mathbb{E}_t\hat{\kappa}_{nit'}\hat{z}_{mt'}\right), \\
\mathbb{E}_t\hat{w}_{nt'}+\mathbb{E}_t\hat{l}_{nt'} &= \sum_i\frac{\bar{\lambda}_{int'}\bar{w}_{it'}\bar{l}_{it'}}{\bar{w}_{nt'}\bar{l}_{nt'}}\left(\mathbb{E}_t\hat{\lambda}_{int'}+\mathbb{E}_t\hat{w}_{it'}+\mathbb{E}_t\hat{l}_{it'}\right) \\
&+\frac{1}{2}\sum_{i,m}\frac{\bar{\lambda}_{int'}\bar{w}_{it'}\bar{l}_{it'}}{\bar{w}_{nt'}\bar{l}_{nt'}}\left[\mathbf{1}(m=i)-\frac{\bar{\lambda}_{mnt'}\bar{w}_{mt'}\bar{l}_{mt'}}{\bar{w}_{nt'}\bar{l}_{nt'}}\right] \\
&\Big(\mathbb{E}_t\hat{w}_{it'}\hat{w}_{mt'}+\mathbb{E}_t\hat{\lambda}_{int'}\hat{\lambda}_{mnt'}+\mathbb{E}_t\hat{l}_{it'}\hat{l}_{mt'}+2(\mathbb{E}_t\hat{\lambda}_{int'}\hat{w}_{mt'}+\mathbb{E}_t\hat{\lambda}_{int'}\hat{l}_{mt'}+\mathbb{E}_t\hat{w}_{it'}\hat{l}_{mt'})\Big).
\end{aligned}
$$

where $\bar{\psi}_{int'+1}=\frac{\bar{\mu}_{int'}\bar{l}_{it'}}{\bar{l}_{nt'+1}}$. The equilibrium objects are $\left\{\mathbb{E}_t\hat{v}_{it'},\mathbb{E}_t\hat{\mu}_{nit'},\mathbb{E}_t\hat{l}_{nt'},\mathbb{E}_t\hat{w}_{it'},\mathbb{E}_t\hat{P}_{it'}\right\}_{i=1,n=1,t'=t}^{N,N,T}$

[4] When t'=t, the expectation operators in the period-t outcomes are redundant. For example, $\mathbb{E}_t\hat{P}_{nt}=\hat{P}_{nt}$.

and the additional inputs compared to Proposition 1 are the covariance between these variables, e.g., $\{\mathbb{E}_t \hat{v}_{it'} \hat{v}_{nt'}, \mathbb{E}_t \hat{w}_{it'} \hat{w}_{nt'}, \mathbb{E}_t \hat{P}_{it'} \hat{P}_{nt'}\}$.

Algorithm. We develop the following simulation-based algorithm.

i. Start from period $t = t_0$, simulate S paths of $z_{t_0}^T \equiv (z_{t_0}, z_{t_0+1}, z_{t_0+2}, ...z_T)$ from agents' belief distribution in t_0.

ii. Solve the first-order approximation for each $s = 1, .., S$ using Proposition 1. Denote the solution $\hat{x}_{t'}(s)$ for all $t' \geq t_0$, where $\hat{x} \in \{\hat{w}_{nt'}, \hat{P}_{nt'}, \hat{v}_{t'+1}, etc.\}_{t'=t_0}^T$.

iii. Approximate $\mathbb{E}_{t_0} \hat{x}_{t'} \hat{y}_{t'} \approx \frac{1}{S} \sum_{s=1}^S \hat{x}_{t'}(s) \hat{y}_{t'}(s)$ and then plug $\mathbb{E}_{t_0} \hat{x}_{t'} \hat{y}_{t'}$ into the system of equations (A.4) and solve for the remaining first-order terms, which gives the decision in period t_0 and the expected outcomes for $t' > t_0$ based on the belief in t_0.

iv. Move the economy to $t = t_0 + 1$, repeat the above process until $t = T$.

A.4 Extension to Incorporate Heterogeneous Beliefs

Without loss of generality, assume that there are two different beliefs about the future path of fundamental productivities, each held by one group of agents.[5] We denote the two groups of agents by A and B, and their beliefs by $f^g(z_{t+1}|z^t)$, $g \in \{A, B\}$. We assume that the existence of these heterogeneous beliefs is common knowledge. In other words, agents in one group are aware of the existence of agents in the other group with a different belief but think that their own belief is the correct one. For simplicity, we assume that agents' belief types do not change, but it is feasible to allow belief types to evolve stochastically.[6] As all decisions of all agents interact in equilibrium, each group makes decisions considering the decisions of the other group guided by their own beliefs.

We now define the stochastic sequential equilibrium in this environment.

Definition 2. ***A stochastic sequential equilibrium with heterogeneous beliefs*** *is a set of state-contingent allocations* $\{\mu_{nit}^g(z^t), \lambda_{nit}(z^t), l_{nt}^g(z^{t-1})\}_{n=1,i=1,t=1,g\epsilon\{A,B\}}^{N,N,T}$, *prices* $\{w_{nt}(z^t), P_{nt}(z^t)\}_{n=1,t=1}^{N,T}$, *and location values* $\{v_{nt}^g(z^t)\}_{n=1,t=1,g\epsilon\{A,B\}}^{N,T}$, *such that* $w_t(z^t), P_t(z^t), \lambda_{nit}(z^t)$ *solve the static trade equilibrium defined by bilateral shares (6), local prices (7), and the labor market clearing condition (8), and for each group* g, $v_{nt}^g(z^t), \mu_{nit}^g(z^t), l_t^g(z^{t-1})$ *solve*

[5] Our approach generalizes to a finite number of belief types, although the computational burden increases with the types of beliefs. Even with just two belief types, however, through the composition of agents' belief types, our model can accommodate variations across locations in average beliefs about fundamentals.

[6] For example, agents' belief type can follow a Markov-switching model. By allowing the switching probability to differ across locations based on local conditions such as the past fundamentals or the current share of each belief type, the model can accommodate endogenous location-specific beliefs.

the dynamic migration decisions given by value function (3), gross flows equation (4), and law of motion of labor (5).

Under Definition 2, agents correctly anticipate the decisions of all other agents in all realizations of z^t but disagree on the likelihood of z^t. In this sense, this definition corresponds to an equilibrium of full rationality. Solving this equilibrium fully is infeasible for the same reason that solving the equilibrium with homogeneous beliefs is infeasible. In fact, with heterogeneous beliefs, even the local approximation solution method established in Proposition 1 involves a curse of dimensionality akin to the one in the literature of higher-order beliefs. In what follows, we first discuss the nature of this problem and then make progress toward tractability by proposing an approximate solution that maintains the generality of the rest of our framework.

Suppose we follow the approach in Proposition 1. Forward-looking decisions by group A depend on the expected future path of real wages, which are shaped by group B's migration decisions, $\mu^B_{nit}(z^t)$. This requires group A to form beliefs about group B's future actions. Since group B's decisions depend on their beliefs about group A's actions, the problem introduces an infinite regress of higher-order beliefs—A's belief about B's belief about A's belief, and so on. This process, illustrated in Figure A.1a, creates an infinitely branching tree of expectations where each future belief state must be evaluated. For example, A's belief about B's decisions depends on A's belief about B's belief about A's decisions further in the future, continuing indefinitely. This implies that a first-order approximation requires solving for an infinite chain of beliefs about future decisions and fundamental productivity (e.g., $\mathbb{E}^A_1\mathbb{E}^B_2\mathbb{E}^A_3 z_4$...), which is computationally infeasible when the state space, z_t, is continuous.

To retain tractability, we impose three alternative structures on heterogeneous beliefs. We emphasize that these assumptions do not restrict agents' beliefs about fundamentals, but their beliefs about other people's future beliefs. These assumptions make the model solvable by collapsing the infinite hierarchy of beliefs.

Assumption 1. *At period t, each type thinks (wishfully) that in t + 1 and beyond, the new data will vindicate their own belief and convince the other type.*

Assumption 2. *Agents maintain their beliefs and think the agents in the other group will also maintain their beliefs; however, group A thinks group B thinks that A will be convinced by the data in the future, and vice versa.*

Assumption 3. *Agents think their beliefs will converge in the future.*

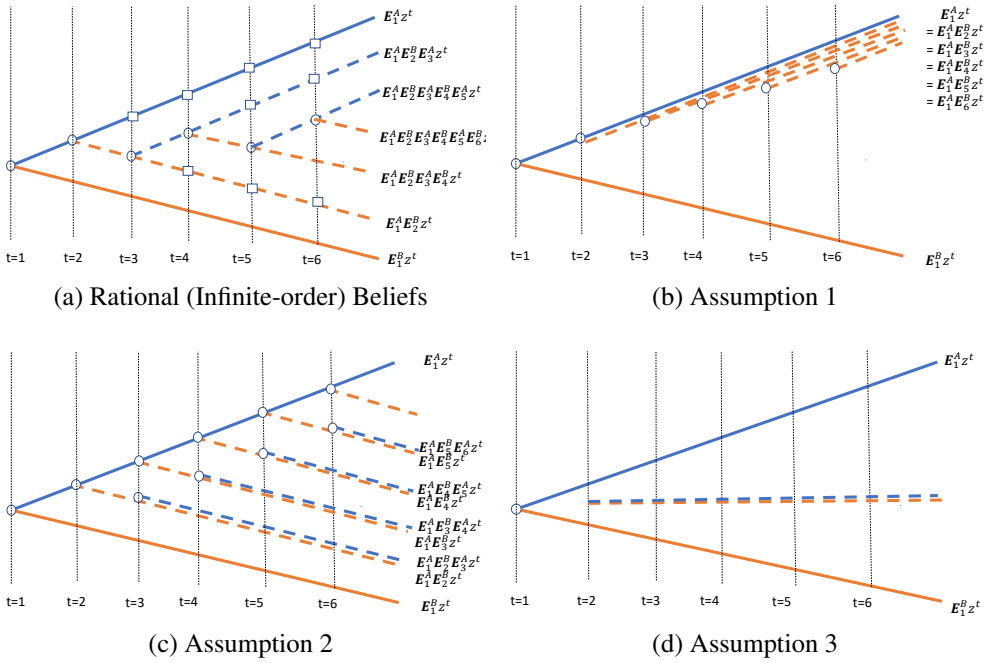

(a) Rational (Infinite-order) Beliefs

(b) Assumption 1

(c) Assumption 2

(d) Assumption 3

Figure A.1: Structures of Heterogeneous Beliefs

Figures A.1b, A.1c, and A.1d illustrate the structures of higher-order beliefs under these three assumptions. Under the first assumption, each group expects that realized fundamental productivities will convince the other group that their own beliefs are correct. This is tractable because it truncates the belief hierarchy at the second order; an agent from group A forecasts the future as if everyone will share A's beliefs from the next period onward, eliminating the need to model a differing belief system for future periods. Under the second assumption, group A thinks that group B thinks group A will be convinced and vice versa, so only up to third-order beliefs appear. Moving from Assumption 1 to Assumption 2, the complexity of the model increases. If we add more and more higher-order beliefs, the model will approach the rational stochastic equilibrium at the expense of computational complexity. Instead of going into even higher order, one can consider a setting (Assumption 3) in which both groups believe that they will converge to a common belief in the next period. These three assumptions capture a wide range of empirically relevant settings.[7] By imposing any of these assumptions, we are able to retain tractability in the stochastic spatial framework with heterogeneous beliefs, as established in Proposition A.1.

Proposition A.1. *Under each of the three assumptions, the local solution to the stochas-*

[7]For example, in lab experiments of coordination games, most subjects engage in less than three levels of reasoning (Mauersberger and Nagel, 2018).

tic spatial model can be characterized by a system of linear equations and can be solved analogously to Proposition 1.

Proof. See Online Supplement C.4. □

Intuitively, under these structures, higher-order beliefs about endogenous outcomes, which are relevant for solving the decision today, can be reduced to lower-order beliefs by the law of iterative expectations. For example, A's belief about B's belief about A's future decisions are simplified to A's belief about A's own future decisions under Assumption 1, and to $A's$ belief about B's future decisions under Assumption 2, thereby simplifying the system of equations that characterizes A's problem. The system of equations that characterizes B's problem can be simplified analogously. The equilibrium is a Nash equilibrium between the decisions of A and B. In Online Supplement C.4, we develop iterative algorithms to solve the problems of A and B jointly at each period, which moves the economy to the subsequent period.

This framework for heterogeneous beliefs can be tractably combined with the second-order accurate solution and the methods for ex-post analysis.

Proposition A.2. *Under Assumptions 1–3, the stochastic spatial framework with heterogeneous beliefs can be solved with second-order accuracy as described in Section 2.6, so uncertainty can be incorporated. The results described in Sections 2.6 and A.4 can also be applied to ex-post studies in which researchers observe outcomes shaped by the decisions of agents whose beliefs are uncertain or heterogeneous, and are interested in finding the allocation corresponding to a counterfactual fundamental path.*

Proof. Online Supplement C.5 lists the full system of equations for these problems. □

The intuition for incorporating uncertainty is as follows. Under the maintained assumptions, each group computes possible future scenarios by simulating draws from a probability distribution over future fundamentals. Under Assumption 1, this distribution is agents' own beliefs—because they anticipate the other group to be convinced in the upcoming period. In this case, the second-order terms can be simulated under a homogeneous-belief setup. Under Assumptions 2 and 3, the subjective distribution for each group is more involved because agents form second- and third-order beliefs. Nevertheless, because the assumptions truncate the high-order beliefs, the simulations are still feasible. Once the second-order terms are calculated from the simulations, we can compute agents' decisions with second-order accuracy analogously to Proposition A.1.

This framework can be applied to ex-post studies, where a backward-induction algorithm on historical allocations can recover the sequence of uncertain or heterogeneous beliefs that best rationalize the data, which then serve as the basis for counterfactuals. We provide algorithms for such exercises in Online Supplement C.5. There, for heterogeneous beliefs in ex-post counterfactuals, our approach takes into account whether the migration flows *by belief type* are observed, focusing on two polar cases: one in which the researcher can observe migration flows by belief type in all periods, and one in which the researcher cannot observe migration flows by belief type but can observe total bilateral migration for each period and one cross-sectional distribution of population by belief type. We show how to conduct counterfactuals in both cases and discuss how to proceed when the data available to the researcher fall between the two cases.

A.5 Accuracy of Second-Order Approximations

We illustrate the quality of second-order approximations by comparing it with the non-linear solution in a deterministic model. We consider an economy with five symmetric regions. We simulate a one-time shock that increases Region 5's fundamental productivity z_{5t} by 100%, while keeping all other fundamentals constant.

Because of the shock, workers move from other regions to region 5. Figure A.2a plots the change in region 5 and other regions' population in the new steady state. Region 5's population almost doubles; population in other regions decreased substantially. The figure shows that despite the large shock, the second-order solution almost exactly matches the non-linear solution. Figure A.2b produces the transitional dynamics of region 5 population. The second-order solution tracks the nonlinear model closely.

Appendix B Applications

B.1 Multi-Sector Extension

Our quantitative applications use an extended model with N locations (indexed by n, i, m, o, h) and J sectors (indexed by j, k, s). We describe the general model below; step-by-step algebra on the second-order approximations of this model can be found in the Online Supplement.

Preferences. A worker's real consumption in location-sector nj is their wage deflated by the local price index, $C_t^{nj} = w_t^{nj} / P_t^n$. The price index is a Cobb-Douglas aggregator of

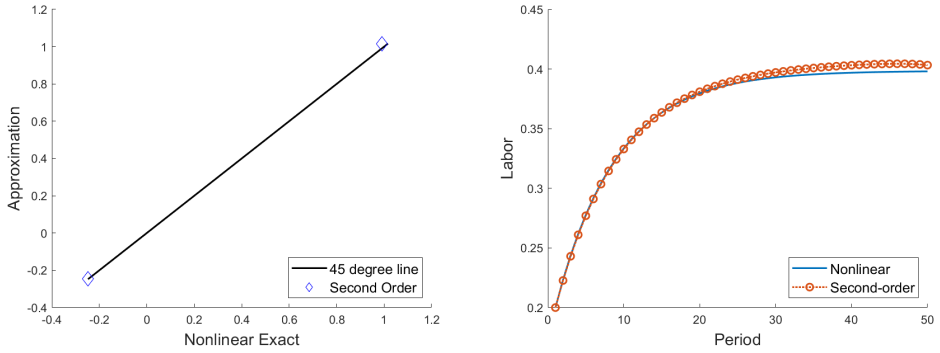

(a) Labor Allocation Across Regions (b) Labor Allocation in a Region with a Shock

Figure A.2: Accuracy of Second-Order Approximation

Note: Panel (a) plots the percentage change in labor allocation for each region under the non-linear model (x-axis) against the second-order approximation (y-axis) at the new steady state. Points closer to the 45° line indicate smaller approximation errors. Panel (b) shows the labor allocation in the region experiencing the positive productivity shock over time under the nonlinear model (solid line) and the second-order approximation (circles).

sectoral prices P_t^{nj} with expenditure shares α^{nj}: $P_t^n \equiv \prod_{j=1}^J (\frac{P_t^{nj}}{\alpha^{nj}})^{\alpha^{nj}}$.

Production, prices, and trade shares. Intermediate goods ω^j are produced using Cobb–Douglas technology combining labor and materials M_t^{nj} in sector j, scaled by efficiency term $T_t^{nj}(\omega^j)$:

$$q_t^{nj}(\omega^j) = T(\omega^j)(l_t^{nj}(\omega^j))^{\gamma^{nj}}(M_t^{nj}(\omega^j))^{1-\gamma^{nj}}.$$

These intermediate goods are aggregated into sectoral final goods. The resulting sectoral price index in nj and its implied share of j goods that n purchases from i are given by input costs, productivity, and trade costs across all source locations:

$$P_t^{nj}(z^t) \propto \left[\sum_{i=1}^N z_t^{ij}\left(c_t^{ij}(z^t)\kappa_t^{nj,ij}\right)^{-\theta^j}\right]^{-1/\theta^j}, \qquad \lambda_t^{nj,ij}(z^t) = \frac{z_t^{ij}\left(c_t^{ij}(z^t)\kappa_t^{nj,ij}\right)^{-\theta^j}}{\sum_{h=1}^N z_t^{hj}\left(c_t^{hj}(z^t)\kappa_t^{nj,hj}\right)^{-\theta^j}},$$

where $c_t^{ij}(z^t) = (\gamma^{ij})^{-\gamma^{ij}}(1-\gamma^{ij})^{-(1-\gamma^{ij})}(w_t^{ij}(z^t))^{\gamma^{ij}}(P_t^{ij}(z^t))^{1-\gamma^{ij}}$.

Market clearing condition. Let X_t^{ij} denote total expenditure on sector-j goods in region i at time t. Goods and labor market clearing then require:

$$X_t^{ij}(z^t) = (1-\gamma^{ij})\sum_{n=1}^N \lambda_t^{nj,ij}(z^t)X_t^{nj}(z^t) + \alpha^{ij}\sum_{k=1}^J w_t^{ik}(z^t)l_t^{ik}(z^t),$$
$$w_t^{nj}(z^t)l_t^{nj}(z^{t-1}) = \gamma^{nj}\sum_{i=1}^N \lambda_t^{ij,nj}(z^t)X_t^{ij}(z^t).$$

Migration decision. Workers choose a location-sector ik to maximize lifetime utility, facing migration costs $m_t^{nj,ik}$. The value function is:

$$v_t^{nj}(z^t) = U(C_t^{nj}(z^t)) + \nu \log[\sum_{i=1}^{N}\sum_{k=1}^{J} \exp(\beta v_{t+1}^{ik}(z^t) - m_t^{nj,ik})^{1/\nu}].$$

The migration decision and the law of motion for labor is:

$$\mu_t^{nj,ik}(z^t) = \frac{\exp[(\beta v_{t+1}^{ik}(z^t) - m_t^{nj,ik})^{1/\nu}]}{\sum_{m=1}^{N}\sum_{h=1}^{J} \exp[(\beta v_{t+1}^{mh}(z^t) - m_t^{nj,mh})^{1/\nu}]}, \qquad l_{t+1}^{nj}(z^t) = \sum_{i=1}^{N}\sum_{k=1}^{J} \mu_t^{ik,nj}(z^t)\, l_t^{ik}(z^{t-1}).$$

The Online Supplement defines the equilibrium for the multi-sector model with intermediate inputs (Section D.1) and under heterogeneous beliefs (Section D.2). We provide the corresponding first- and second-order derivatives used in the quantitative applications (Sections D.3 and D.4) and present the algorithm for solving the model (Section D.5).

B.2 China Shock Application

B.2.1 Geographic and Time Coverage

List of countries and regions. We calibrate the model to the 50 U.S. states and 28 other countries/regions including a synthetic rest of the world. These countries/regions include Australia, Austria, Balkans (aggregating Bulgaria, Croatia, and Greece), Baltic States (aggregating Estonia, Latvia, and Lithuania), Benelux (Belgium, Netherlands, Luxembourg), Brazil, Canada, Switzerland, China, Central Europe (aggregating Poland, Czech Republic, Slovakia, Slovenia, and Hungary), Germany, Spain, France, United Kingdom, Indonesia, India, Ireland, Italy, Japan, South Korea, Mexico, Nordic Countries (aggregating Denmark, Finland, Norway, and Sweden), Portugal, Romania, Russia, Turkey, Taiwan, and the rest of the world. Among the 50 U.S. states, the District of Columbia is aggregated into Virginia.

Time horizon. Our analysis covers 2000–2099. Each period is treated as a year and we set $\beta = 0.96$.

B.2.2 Data Construction

In this application, we treat 2000–2016 as the in-sample period for which we construct the evolution of the economy from the data. We treat years after 2016 as out of the sample.

Thus, we need to construct yearly data between 2000 and 2016. This subsection explains the sources of information and data construction in detail.

Sector aggregation. We incorporate two broad sectors, manufacturing and non-manufacturing, and aggregate sector-level information from various sources of data to these two broad sectors, as described in Table B.2.

Employment and migration. We construct the initial allocation of labor and the subsequent migration shares as follows:

1. **Initial labor allocation l_0.** We construct the initial labor allocation in 2000 (l_0^{nj}) using the employment data across sectors and states obtained from the BEA. We normalize the level so that $\sum_{n=1}^{R}\sum_{j=1}^{J} l_0^{nj} = 1$.
2. **Migration share.** We construct annual migration shares for 2000–2014 using job-to-job flow data from the Longitudinal Employer-Household Dynamics Database (LEHD), with the residual share in each state-sector cell treated as stayers.[8] Since the LEHD data are incomplete for some states between 2000–2009 (see Table B.1), we impute the missing values using migration flows from Caliendo et al. (2019).[9]

Table B.1: States Not Covered by LEHD

Year	States
2000	AL, AZ, AR, KY, MA, MS, NH, WY
2001-2002	AZ, AR, MA, MS, NH, MI
2003	AR, MA
2004-2009	MA

We obtain value added, and input and final consumption shares across sectors and countries from the World Input-Output Database (WIOD), using data from 2007—the midpoint of the sample period. We obtain gross output and trade flows across sectors and countries from the Eora Database.[10] Using this information, we construct the following variables:

Sectoral shares in final consumption. To calculate sectoral shares in final consump-

[8] While LEHD J2JOD data contain information on the stayers, we do not use it since it does not count those who did not change their job.

[9] Caliendo et al. (2019) construct migration flows between states and sectors over this period using data from the American Community Survey. Their data cover 23 sectors, including a non-employment sector, which we remove. We then normalize the remaining flows to match our data sequence and aggregate them into our two sectors. Specifically, we map sectors 2-13 in their data to manufacturing and sectors 14-23 to non-manufacturing.

[10] We use this series to measure the productivity of China and other countries before the year 2000 through model inversion. We use Eora instead of WIOD for this purpose because the latter does not extend long enough before 2000 for estimating productivity catch up of China.

tion, α^{nj}, we sum the final expenditures on sector j goods by households, nonprofit organizations, and governments in each country n and divide the sum by the total final absorption of n country in the WIOD. We do this for both foreign countries and the United States, assuming that α^{nj} is common across U.S. states.

Value-added share in intermediate goods production. We compute the share of value added in the gross output γ^{nj} and the intermediate consumption share $1 - \gamma^{nj}$ by dividing total intermediate input uses in sector j of country n by the gross output of the sector-country cell. Again, γ^{nj} are identical across U.S. states but differ by country.

Production and absorption. We obtain national-level production and absorption data directly from Eora. To construct these values for individual U.S. states, we combine state-level value-added data from the U.S. Bureau of Economic Analysis (BEA) with national-level sectoral consumption shares (α^{nj}) and value-added shares (γ^{nj}). We then combine these data with the model's market clearing conditions to calculate state-specific production and absorption, which we scale to match the aggregate statistics from the Eora.

International and inter-state trade. Trade shares between non-U.S. countries are directly computed using the Eora. Since the model is at the state-sector level within the U.S., we need to construct inter-state trade and trade between U.S. states and other countries. We do this by supplementing the Eora with additional information, as described below:

i. Trade between U.S. states and other countries

   We use each state's share in U.S. output in each sector as the state's share in U.S. exports to each country and imports from each country. We use output instead of imports and exports to calculate the share because the U.S. trade online database does not cover the earlier part of our sample.

ii. Trade between U.S. states

   (a) Manufacturing: The CFS is available for years 2002, 2007, 2012, and 2017. We linearly interpolate the trade shares calculated from the CFS for all years between 2002 and 2016.[11] We multiply these shares by the level of manufacturing-sector domestic sales in the United States according to the WIOD for each year to obtain the level of trade flows between U.S. states.

   (b) Non-manufacturing: The non-manufacturing sector includes both agriculture and services, so we assume that it is tradable. We first estimate the elasticities based on a gravity equation using the inter-state trade flow in manufacturing constructed above. Using these estimates and non-manufacturing absorption and

[11] For 2000 and 2001, we assume the trade shares are the same as in 2002.

production, we then impute the trade flows in the non-manufacturing sector. The gravity specification we use is

$$\ln(X_t^{in}) = \beta_1 \ln(X_t^n) + \beta_2 \ln(Y_t^i) + \beta_3 \ln(\text{dist}^{in}) + \beta_4 \mathbb{1}(i = n), \tag{B.1}$$

where X_t^{in} is the bilateral flow between i and n, $\ln(X_t^n)$ is the absorption in the destination, and Y_t^i is the production in the source country. Absorption and production are constructed as described above. Bilateral distance is aggregated from the county-pair level, weighted by population. We scale the level of transactions to match the domestic sales in the U.S. non-manufacturing sector from the WIOD. The estimated coefficients using the 2007 data are $\hat{\beta}_1 = 0.650$, $\hat{\beta}_2 = 0.776$, $\hat{\beta}_3 = -1.347$, and $\hat{\beta}_4 = 1.376$.

Measuring productivity catch-up. To estimate China's productivity catch-up process, we measure the productivity of the United States and China from the data. In our model, the productivity parameter of country n in sector j, z_t^{nj} is

$$z_t^{nj} = \lambda_t^{nj,nj} \left(\frac{w_t^{nj}}{P_t^{nj}} \right)^{\theta \gamma^{nj}}.$$

We use this equation to measure z_t^{nj}. The construction of γ^{nj} is as described before. We construct the trade share sequence $\lambda_t^{nj,nj}$ from the Eora database because the WIOD does not cover the earlier part of the sample.[12] For wages and prices, we use Penn World Table (Version 9.1). We measure w_t^{nj} as the ratio between total labor compensation and total number of employment (variable emp in PWT), where labor compensation is calculated as the product of labor share (variable $labsh$) and GDP (variable $cgdpo$). We proxy for the index for manufacturing price using the average of the import and export prices of country n in period t (variables pl_m and pl_x, respectively).

B.2.3 China's Manufacturing Productivity Process and Beliefs

Figure B.1a plots the actual productivity process for China's manufacturing (blue line) against three alternative growth scenarios for China: 5.5%, 3.5%, and 2.6% per annum. Our main counterfactual uses the 3.5% growth rate, reaching the peak in 2060.

Figure B.1b shows the recovered beliefs about China's manufacturing productivity. Ini-

[12] Our quantitative analysis focuses on after 2000, which WIOD covers, however, we also show the evolution of the relative productivity between the U.S. and China since the early 1990s, which WIOD does not cover.

Table B.2: List of Sectors and Datasets

Dataset	Agriculture	Manufacturing	Service
Eora	Agriculture, Fishing, Mining	Manufacturing	The rest
WIOD	Agriculture, Forestry, Fishing, Mining	Manufacturing	The rest
BEA:GDP	Agriculture, forestry, fishing, and hunting Mining, quarrying, and oil and gas extraction	Manufacturing	The rest
BEA:EMPLOYMENT	Farm employment (linecode 70) Forestry, fishing, and related activities Mining, quarrying, and oil and gas extraction	Manufacturing (linecode 500)	The rest
CENSUS (US TRADE ONLINE)	111 Agricultural Products 112 Livestock & Livestock Products 113 Forestry Products, Nesoi 114 Fish, & Other Marine Products 211 Oil & Gas 212 Minerals & Ores	The rest	N/A
CFS	N/A	Manufacturing	N/A
LEHD	Agriculture, forestry, fishing, and hunting Mining, quarrying, and oil and gas extraction	Manufacturing (Sectors 31–33)	The rest

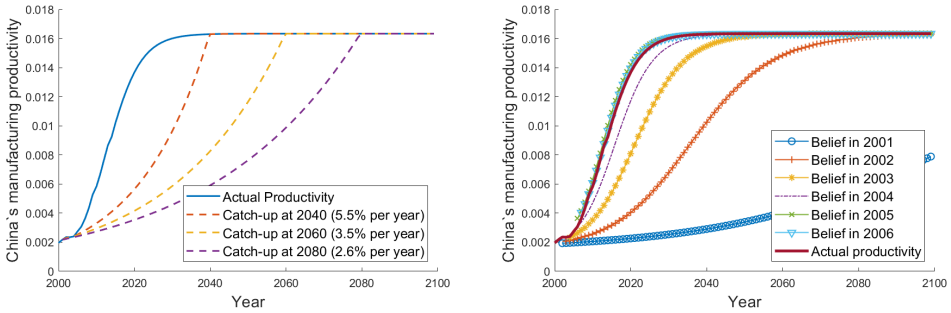

(a) China's Manufacturing Productivity (b) Belief About China's Productivity

Figure B.1: China's Manufacturing Productivity Process and Beliefs

Note: Panel (a) presents the actual productivity process of China's manufacturing productivity (blue solid line) and three counterfactual productivity processes corresponding to different growth rates (dashed lines). Panel (b) plots the actual productivity process of China's productivity (red solid line) and the recovered beliefs about this process formed at different years (all others).

tially pessimistic about China's productivity, agents' beliefs gradually revised upwards, aligning with the actual productivity (solid line) by 2006.

B.2.4 Employment Effect of the China Shock

Figure B.2a shows the long-run employment effects of the China shock until 2099. It compares the decline in manufacturing employment in our model with evolving beliefs (blue line) to a perfect foresight model (red line). In both models, the decline intensifies even after 2016 (the end of sample) as the difference between the actual (solid line in Figure B.1a) and counterfactual productivity paths of China (3.5% growth scenario in Figure B.1a) continue to increase over time. The employment decline eventually reverses as both the actual

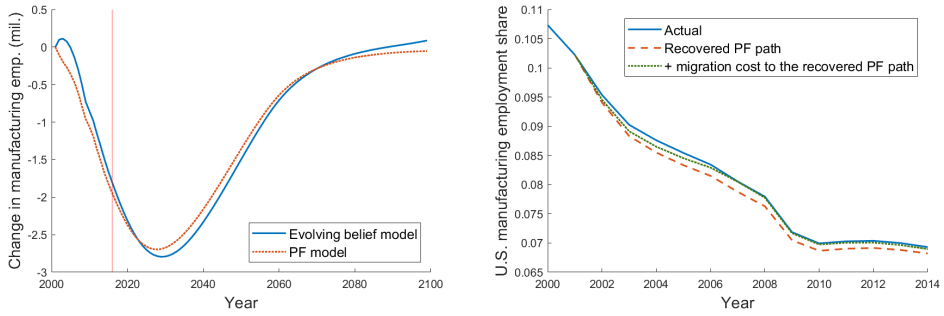

(a) Employment Effect of the China Shock (b) Employment Share Under Higher Migration Cost

Figure B.2: Employment Effects of the China Shock and the Role of Migration Cost

Note: Panel (a) presents the employment effects of the China shock, measured as the difference between the economy with the shock and a counterfactual without it, under both evolving beliefs (blue solid line) and perfect foresight (red dashed line). Panel (b) shows that the actual manufacturing employment path (blue solid line) and the path inferred under perfect foresight (red dashed line) can be closed by adding a 50% migration cost to the perfect foresight model (green dotted line), quantifying the impact of evolving beliefs.

and the counterfactual productivity converge to the same level.

B.2.5 Assessing the Importance of Evolving Beliefs

We assess the importance of accounting for evolving beliefs in terms of migration costs. The red dashed line in Figure B.2b plots the perfect foresight path recovered from our model with evolving belief—the outcome if agents had perfect foresight in 2001. This path implies a faster reallocation out of manufacturing than the actual data (blue solid line) because agents anticipated China's productivity catch-up and exited the sector sooner.

A researcher assuming perfect foresight would treat the actual data (blue solid line) as the perfect foresight path. The difference between this and our recovered path (red dashed line) therefore reflects the importance of accounting for evolving beliefs, a gap that reached 0.2 percentage points of employment (or 0.3 million jobs) by 2007.

To quantify the importance of evolving beliefs, we conduct a counterfactual experiment starting from the red dashed path, increasing the cost of migrating from manufacturing by a factor of (1+x) over 2001–2007. We choose x so that this counterfactual path has the same manufacturing employment share as the data in 2007. We find that when x=50%, this counterfactual path (the green dotted line) matches the 2007 share, implying that a perfect foresight model would need an approximately 50% higher migration cost to account for the observed path.

B.3 Climate Change Application

Most of the data we use for the climate change application are the same as those used in the China shock application, so the same data sources and procedures apply. Below, we explain the additional data construction for the climate change application.

Regions. We focus on the same set of countries and (aggregated) regions as in the China shock application, explained in Section B.2.

Definition of the five broad U.S. regions. We report some of the results for five broad U.S. regions: Northeast (NE), Southeast (SE), Midwest (ME), Southwest (SW), and West (WE). These regions are defined as follows:

Northeast: Connecticut, Delaware, Maine, Maryland, Massachusetts, New Hampshire, New Jersey, New York, Pennsylvania, Rhode Island, and Vermont.

Southeast: Alabama, Arkansas, Florida, Georgia, Kentucky, Louisiana, Mississippi, North Carolina, Tennessee, South Carolina, Virginia, and West Virginia.

Midwest: Iowa, Illinois, Indiana, Kansas, Michigan, Missouri, Ohio, Wisconsin, Minnesota, Nebraska, North Dakota, and South Dakota.

Southwest: Arizona, New Mexico, Oklahoma, and Texas.

West: Alaska, California, Colorado, Hawaii, Idaho, Montana, Nevada, Oregon, Utah, Washington, and Wyoming.

Time horizon. We treat 2013–2014 as the starting point of the model and collect from various data sources the information necessary for constructing this initial economy. Our analysis focuses on the years 2014–2100. We treat each period in the model as two years. Consistent with this treatment, we set $\beta = 0.92$, corresponding to an annual discount rate of 4%; we also use two-year migration flows, constructed based on the job flow information provided by the Longitudinal Employment History Database, described below.

B.3.1 Data Construction

To construct the initial economy, we need information on migration, trade, and production around 2014. To solve for the evolution of the economy over 2014–2100, we also need temperature projection (mean and variance) and the distribution of climate skeptics. This subsection explains the sources of information and data construction in detail.

Sector aggregation. We incorporate three broad sectors—agriculture, manufacturing, and service—and aggregate sector-level information from various sources of data to these three broad sectors. Table B.2 summarizes the aggregation.

Initial migration share. As in the China shock application, we utilize the LEHD J2JOD data to construct the migration shares in 2013. In the climate change application, each period is two years. We compound the one-year migration shares to construct a two-year migration share matrix.

Trade and production data. We obtain production and value-added data from the BEA for the U.S. and from WIOD for other countries.

International and inter-state trade. Trade shares between non-U.S. countries are from the WIOD.[13] We construct inter-state trade and trade between U.S. states and other countries by supplementing the WIOD with additional information, as described below:

i. Trade between U.S. states and other countries
   We use the U.S. trade data provided by the U.S. Census to compute each state's share in U.S. exports to each destination country and U.S. imports from each source country. We then apply these shares to the bilateral trade flows between the United States and each other country in the WIOD to obtain the level of trade flows between U.S. states and other countries.

ii. Trade flows between U.S. states
   (a) Manufacturing: We use 2012 Commodity Flow Survey (CFS) data to obtain the bilateral trade flows across individual U.S. states for the manufacturing sector. We scale the total level of transactions in CFS to match the domestic sales of the manufacturing sector in the United States from the WIOD.
   (b) Agriculture: We estimate the elasticities from the gravity equation (B.1), using the CFS to impute the trade flows as in Section B.2.2. The estimated coefficients are $\hat{\beta}_1 = 0.786$, $\hat{\beta}_2 = 0.745$, $\hat{\beta}_3 = -1.668$, and $\hat{\beta}_4 = 0.848$.
   (c) Service: We assume there is no trade in service between U.S. states.

Using these flows between and within U.S. states, between U.S. states and foreign countries, and between foreign countries constructed, we construct bilateral trade shares.

Temperature projections. We construct temperature projections based on the Climate Impact Lab, accessed on September 9, 2022.[14] The Climate Impact Lab provides multiple carbon emissions scenarios (e.g. RCP 4.5), with three temperature paths under each scenario (i.e., median, 5th percentile, and 95th percentile). In each case, the projection is reported for four broad periods—1986–2005 (historic), 2020–2039, 2040–2059, and 2080–2099—and

[13] We assume that service is non-tradable so we attribute all service consumption expenditures to local production.

[14] Climate Impact Lab updated the projection scenarios in 2024; our raw data are available upon request.

is available for both foreign countries and U.S. states. Our quantification focuses on the RCP 4.5 scenario, which is described by the IPCC as an intermediate scenario with carbon emissions peaking in 2040 and then declining.

Interpolating to bi-yearly frequency. We interpolate the temperature within these time frames linearly to obtain yearly data for the median, 5^{th}, and 95^{th} percentiles in temperature projection. To be specific, we assume the projected temperature in each period is for the middle year of each period (2030, 2050, and 2090). We then linearly interpolate the temperature between those years and 1996 (the middle year between 1986 and 2005).[15] We treat the value of the median path in 2014 from this interpolation as the temperature for 2014.

Variance profile. We adopt the RCP 4.5 median scenario as the average temperature path and use the median and 5^{th} percentile likelihood values to calibrate $(\sigma_{nt}^1)^2$, the variance in the average annual temperature in period t from the perspective of the agent in the first period.[16] Specifically, we choose $(\sigma_{nt}^1)^2$ so that the difference between the median and the 5^{th} percentile temperature projections in t is the same as implied by $(\sigma_{nt}^1)^2$ under the Gaussian assumption on the temperature path. That is,

$$(\sigma_{nt}^1)^2 = \left((\bar{\mathcal{C}}_{nt}^1 - \mathcal{C}_{nt}^{1,5th})/1.645\right)^2,$$

where $\mathcal{C}_{nt}^{1,5th}$ denotes the 5^{th} percentile projection for the temperature in region n for period t from today, and the value 1.645 is the z-score corresponding to the 5th percentile in a normal distribution. We assume that the variance profile depends only on the duration of the forecast. That is, the uncertainty perceived in the forecast made in 1996 about 2010 is the same as the uncertainty perceived in 2030 about 2044. Formally, the variance in the annual average temperature in period t' for an agent in period $t < t'$, denoted by $\sigma_{nt'}^t$, is simply $\sigma_{nt'-t+1}^1$: for example, $\sigma_{n7}^5 = \sigma_{n3}^1$.

B.3.2 Productivity Path for U.S. Regions

Figure B.3 plots the population-weighted average of the change in $\ln(z_t^{ni})$ for U.S. regions relative to 2014 values for each sector calculated according to (19). The panels show that productivity decreases substantially in all three sectors in the South, but the decrease is

[15] For 2090-2100, we extrapolate from the 2051-2090 path.

[16] Note that if we adopt RCP 8.5, which assumes high future emissions and a dramatic expansion of coal use, the uncertainty in the later period would be larger. Similarly, as the 95^{th} percentile scenario implies larger increase in temperature, using the difference between 5^{th} percentile and the median path implies smaller uncertainty. In this sense, our quantification is a conservative estimate of the impact of uncertainty.

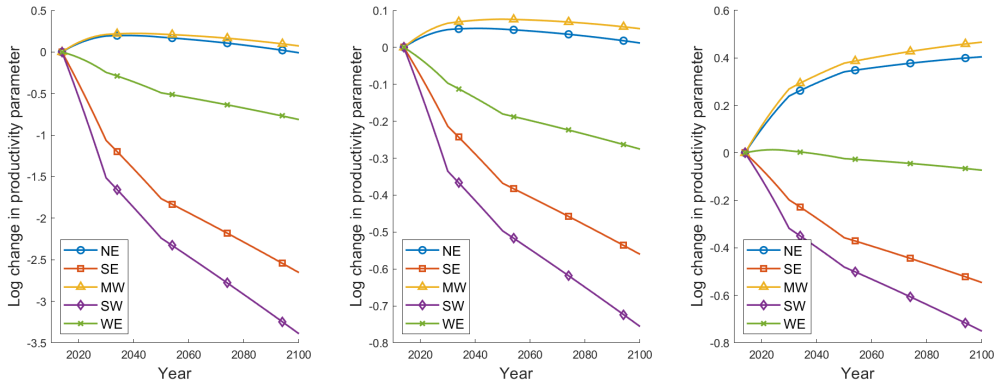

(a) Agriculture (b) Manufacturing (c) Service

Figure B.3: Temperature Projection and Productivity Path

Note: Panels (a), (b), and (c) present the population-weighted change in the productivity parameters for broad U.S. regions under the median temperature path.

particularly pronounced in agriculture, where $\ln(z_t^{nj})$ decreases by almost 300 log points. In terms of the measured TFP, which is given by $(z_t^{nj}/\lambda_t^{nj,nj})^{\frac{1}{\theta}}$, the magnitude of these shocks are scaled down by the trade elasticity and trade selection.

On the other hand, the West sees only a modest decrease in productivity in each sector; the Northwest and Midwest see an increase in productivity. Although not shown in the figure, similar patterns hold for other countries; in general, high-latitude countries such as Canada and Russia experience an increase in productivity from climate change, whereas low-latitude countries such as India suffer productivity loss.

B.3.3 Spatial Heterogeneity in Beliefs about Climate Change

We use the 2014 Yale Climate Opinion Survey to calibrate the initial share of climate skeptics in each state, defined as those answering "No" or "Don't know" to the following question: "What do you think: Do you think that global warming is happening?" We define those answering "yes" as believers of climate change. The average share of believers across U.S. states in 2014 is 62.6% and the standard deviation across states is 5.32%.

Figure B.4 plots the geographic variations in this measure. There is a wide geographic divide, with residents on the coasts generally more inclined to believe climate change.

B.3.4 Welfare under Heterogeneous Beliefs

Figure B.5 plots the welfare effects for believers in individual states.

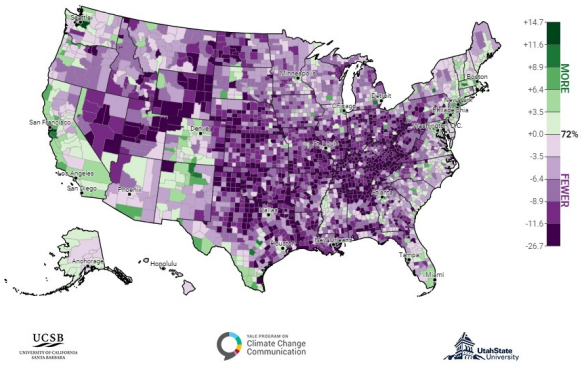

Figure B.4: Share of Adults Who Think Global Warming is Happening

Note: This figure presents the geographical distribution of difference from the national average (72%) in the share of adults who think global warming is happening in 2021. Source: Yale Climate Opinion Map (2021).

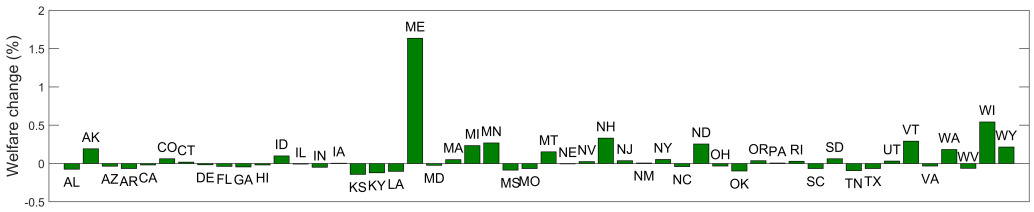

Figure B.5: Welfare of Believers Compared to Perfect Foresight

Note: This figure plots the effect of introducing climate skeptics on the welfare (in consumption equivalents) of the believers of climate change based on the current state of believers, compared to the effect when all agents perfectly anticipate climate change.

B.3.5 Robustness Check

In this subsection, we present results using three different parameter variations. In the first case, we adjust the risk aversion parameter σ; in the second case, we change the size of the productivity shock by half; in the third case, we modify the belief convergence parameter ρ.

Lower risk aversion parameter. Panel A in Table B.3 reports the aggregate US welfare impact of a perfectly anticipated climate change shock ('PF shock') and the additional impact of uncertainty, tested under different risk aversion parameters (σ). As we lower σ from 3 to 2, the results show that lower risk aversion reduces the welfare losses from both the anticipated shock and from the uncertainty.[17]

Smaller productivity shock. As a sensitivity analysis, the last column of Panel A in Table B.3 also reports the welfare change when the productivity shock due to climate change

[17] To see why the welfare impact of the PF shock relative to the no shock increases with σ, note that the change in welfare is measured in consumption equivalents, and is a function of σ: $\delta^{nj} = \left(\frac{(1-\sigma)\left(V_0^{nj} - V_0^{\prime nj}\right)}{\sum_{s=0}^{T} \beta^s \left(C_s^{\prime nj}\right)^{1-\sigma}} + 1 \right)^{\frac{1}{1-\sigma}}$, where δ^{nj} is the welfare change in region-sector nj.

is smaller than in the benchmark analysis. Specifically, we reduce the estimates for the coefficients of the damage function by half (and their standard errors by a factor of $\sqrt{2}$) in Table 1. Halving the magnitude of the productivity shock reduces the aggregate welfare loss from both the perfectly anticipated climate change shock and the associated uncertainty. Nevertheless, the additional loss from uncertainty is still in the same order as the loss from perfectly anticipated climate change.

Higher convergence speed. Panel B in Table B.3 reports the welfare change from climate change when introducing uncertainty, compared to a perfectly foresighted climate change scenario, under different values of the convergence speed parameter ρ in equation (20). A smaller ρ implies faster convergence to the median forecast of temperature. Our findings indicate that a lower ρ reduces the impact of uncertainty on welfare, as faster convergence brings agents' beliefs closer to perfect foresight.

Table B.3: U.S. Welfare Impact of Climate Change under Alternative Specifications

Panel A: Lower Risk Aversion and Productivity Shock				
	$\sigma = 3$	$\sigma = 2.5$	$\sigma = 2$	Half productivity impact ($\sigma = 3$)
No shock - PF shock	-1.60%	-1.36%	-1.31%	-0.68%
PF shock - uncertainty	-1.48%	-1.21%	-0.84%	-0.75%
Panel B: Higher Convergence Speed				
	ρ=0.9	ρ=0.8	ρ=0.5	
Welfare change (%)	-1.481%	-1.479%	-1.294%	